

Lectures on Holomorphic Curves in Symplectic and Contact Geometry

(Work in progress—Version 3.1)

Chris Wendl

INSTITUT FÜR MATHEMATIK, HUMBOLDT-UNIVERSITÄT ZU BERLIN
E-mail address: `wendl@math.hu-berlin.de`

Contents

| | |
|--|-----|
| Preface | v |
| Chapter 1. Introduction | 1 |
| 1.1. Warm up: Holomorphic curves in \mathbb{C}^n | 1 |
| 1.2. Hamiltonian systems and symplectic manifolds | 3 |
| 1.3. Some favorite examples | 5 |
| 1.4. Darboux's theorem and the Moser deformation trick | 6 |
| 1.5. From symplectic geometry to symplectic topology | 10 |
| 1.6. Contact geometry and the Weinstein conjecture | 13 |
| 1.7. Symplectic fillings of contact manifolds | 19 |
| Chapter 2. Local properties | 23 |
| 2.1. Almost complex manifolds and J -holomorphic curves | 23 |
| 2.2. Compatible and tamed almost complex structures | 27 |
| 2.3. Linear Cauchy-Riemann type operators | 30 |
| 2.4. The linearization of $\bar{\partial}_J$ and critical points | 32 |
| 2.5. Linear elliptic regularity | 36 |
| 2.6. Local existence of holomorphic sections | 43 |
| 2.7. Similarity and unique continuation | 45 |
| 2.8. Nonlinear regularity | 51 |
| 2.9. Some tools of global analysis | 56 |
| 2.10. Local existence of J -holomorphic curves | 61 |
| 2.11. A representation formula for intersections | 67 |
| Chapter 3. Fredholm theory | 79 |
| 3.1. Some Banach spaces and manifolds | 79 |
| 3.2. Formal adjoints | 83 |
| 3.3. The Fredholm property | 89 |
| 3.4. The Riemann-Roch formula and transversality criteria | 91 |
| Chapter 4. Moduli spaces | 97 |
| 4.1. The moduli space of closed J -holomorphic curves | 97 |
| 4.2. Classification of pointed Riemann surfaces | 102 |
| 4.3. Fredholm regularity and the implicit function theorem | 113 |
| 4.4. Transversality for generic J | 119 |
| Chapter 5. Bubbling and nonsqueezing | 129 |
| 5.1. Gromov's nonsqueezing theorem | 129 |
| 5.2. Monotonicity in the integrable case | 137 |

| | |
|-------------------|-----|
| 5.3. Bubbling off | 139 |
| Bibliography | 145 |

Preface

The present book-in-progress began as a set of lecture notes written at a furious pace to accompany a graduate course on holomorphic curves that I taught at ETH Zürich in Spring 2009, and repeated at the Humboldt-Universität zu Berlin in the 2009-10 Winter semester. In both iterations of the course, it quickly became clear that my conceived objectives for the notes were not really attainable within the length of the semester, but the project nonetheless took on a life of its own. I have written these notes with the following specific goals in mind:

- (1) To give a solid but readable presentation of the analytical foundations of closed holomorphic curves from a modern perspective.
- (2) To use the above foundation to explain a few of the classic applications to symplectic topology, such as Gromov’s nonsqueezing theorem [Gro85] and McDuff’s results on rational and ruled symplectic 4-manifolds [McD90].
- (3) To use the aforementioned “modern perspective” to generalize everything as cleanly as possible to the case of punctured holomorphic curves, and then explain some applications to contact geometry such as the Weinstein conjecture [Hof93] and obstructions to symplectic fillings [Wen10b].

The choice of topics covered and their presentation is partly a function of my own preferences, as well as my perception of which gaps in the existing literature seemed most in need of filling. In particular, I have devoted special attention to a few topics that are fundamental but not covered in the standard book on this subject by McDuff and Salamon [MS04], e.g. the structure of Teichmüller space and of the moduli space for holomorphic curves of arbitrary genus, and existence results for local J -holomorphic curves. My choice of applications is biased toward those which I personally find the most beautiful and which admit proofs with a very geometric flavor. For most such results, there are important abstract invariants lurking in the background, but one need not develop them fully in order to understand the proofs, and for that reason I have left out topics such as gluing analysis and Gromov-Witten theory, on which I would in any case have nothing to add to the superb coverage in [MS04]. In order to save space and energy, I have also included nothing about holomorphic curves with boundary, but aimed to make up for this by devoting the last third of the book to punctured holomorphic curves, a topic on which I am not aware of any previously available exposition aimed at graduate students.

My personal attitude toward technical details is essentially that of a non-analyst who finds analysis important: what this means is that I’ve tried very hard to create an accessible presentation that is as complete as possible without boring readers who don’t enjoy analysis for its own sake. In contrast to [MS04], I have not put

the discussion of elliptic regularity in an appendix but rather integrated it into the main exposition, where it is (I hope) less likely to be ignored. On the other hand, I have presented such details in less generality than would theoretically be possible, in most places only as much as seems essential for the geometric applications. One example of this is the discussion in Chapter 2 of a local representation formula that is both *weaker* and *easier to prove* than the famous result of Micallef and White [MW95], but still suffices for crucial applications such as positivity of intersections. If some hardcore analysts find this approach lazy, my hope is that at least as many hardcore topologists may benefit from it.

About the current version. Versions 1 and 2 of these notes were the versions written to accompany the actual lecture course, and version 3 is the revision in book form. Version 3.1 contains approximately half of what I hope to include in the finished product: there is not yet any serious material on contact geometry (only a few main ideas sketched in the introduction), but the development of the technical apparatus for closed holomorphic curves is mostly complete. The main thing still missing from this technical development is Gromov's compactness theorem, though a simple case of it is covered in Chapter 5 in order to prove the nonsqueezing theorem. I hope to add the chapter on Gromov compactness in the next revision, along with further chapters covering the special analytical properties of closed holomorphic curves in dimension four, and applications to symplectic 4-manifolds (notably McDuff's results [McD90]).

Acknowledgments. I'd like to thank a number of people who've contributed useful comments and encouragement on this project, most especially Peter Albers, Joel Fish, Paolo Ghiggini, Janko Latschev, Klaus Mohnke and Dietmar Salamon. Thanks also to Urs Fuchs for careful reading and corrections.

A very large portion of what I know about this subject was originally imparted to me by Helmut Hofer, whose unpublished manuscript with Casim Abbas [AH] has also been an invaluable resource for me. Other invaluable resources worth mentioning include of course [MS04], as well as the expository article [Sik94] by Sikorav.

Request. As should by now be obvious, these notes are work in progress, and as such I welcome comments, questions, suggestions and corrections from anyone making the effort to read them. These may be sent to wendl@math.hu-berlin.de.

CHAPTER 1

Introduction

1.1. Warm up: Holomorphic curves in \mathbb{C}^n

The main subject of these lectures is a certain interplay between *symplectic* structures and *complex* (or rather *almost complex*) structures on smooth manifolds. To illustrate the connection, we consider first the special case of holomorphic curves in \mathbb{C}^n .

If $\mathcal{U} \subset \mathbb{C}^m$ is an open subset and $u : \mathcal{U} \rightarrow \mathbb{C}^n$ is a smooth map, we say that u is **holomorphic** if its partial derivatives $\frac{\partial u}{\partial z_j}$ all exist for $i = j, \dots, m$, i.e. the limits

$$\frac{\partial u}{\partial z_j} = \lim_{h \rightarrow 0} \frac{u(z_1, \dots, z_{j-1}, z_j + h, z_{j+1}, \dots, z_m) - u(z_1, \dots, z_m)}{h}$$

exist, where h is complex. This is the obvious generalization of the notion of an analytic function of one complex variable, and leads to an obvious generalization of the usual Cauchy-Riemann equations.

We will find the following equivalent formulation quite useful. Let us identify $\mathbb{C}^n = \mathbb{R}^{2n}$ by regarding $(z_1, \dots, z_n) \in \mathbb{C}^n$ as the real vector

$$(p_1, q_1, \dots, p_n, q_n) \in \mathbb{R}^{2n},$$

where $z_j = p_j + iq_j$ for $j = 1, \dots, n$. Then at every point $z \in \mathcal{U} \subset \mathbb{C}^m$, our smooth map $u : \mathcal{U} \rightarrow \mathbb{C}^n$ has a differential $du(z) : \mathbb{C}^m \rightarrow \mathbb{C}^n$, which is in general a *real*-linear map $\mathbb{R}^{2m} \rightarrow \mathbb{R}^{2n}$. Observe also that for any number $\lambda \in \mathbb{C}$, the complex scalar multiplication

$$\mathbb{C}^n \rightarrow \mathbb{C}^n : z \mapsto \lambda z$$

defines a real-linear map from \mathbb{R}^{2n} to itself. It turns out that u is holomorphic if and only if its differential at every point is also *complex*-linear: in particular it must satisfy $du(z)\lambda V = \lambda \cdot du(z)V$ for every $V \in \mathbb{C}^m$ and $\lambda \in \mathbb{C}$. Since $du(z)$ is already real-linear, it suffices to check that $du(z)$ behaves appropriately with respect to multiplication by i , i.e.

$$(1.1) \quad du(z) \circ i = i \circ du(z),$$

where we regard multiplication by i as a linear map on \mathbb{R}^{2m} or \mathbb{R}^{2n} .

Exercise 1.1. Show that (1.1) is equivalent to the usual Cauchy-Riemann equations for smooth maps $u : \mathcal{U} \rightarrow \mathbb{C}^n$.

If $m = 1$, so \mathcal{U} is an open subset of \mathbb{C} , we refer to holomorphic maps $u : \mathcal{U} \rightarrow \mathbb{C}^n$ as **holomorphic curves** in \mathbb{C}^n . The choice of wording is slightly unfortunate if you like to think in terms of *real* geometry—after all, the image of u looks more like a

surface than a curve. But we call u a “curve” because, in complex terms, it is a one-dimensional object.

That said, let us think of holomorphic curves for the moment as real 2-dimensional objects and ask a distinctly real 2-dimensional question: what is the area traced out by $u : \mathcal{U} \rightarrow \mathbb{C}^n$? Denote points in \mathcal{U} by $s + it$ and think of u as a function of the two real variables (s, t) , with values in \mathbb{R}^{2n} . In these coordinates, the action of i on vectors in $\mathbb{C} = \mathbb{R}^2$ can be expressed succinctly by the relation

$$i\partial_s = \partial_t.$$

We first have to compute the area of the parallelogram in \mathbb{R}^{2n} spanned by $\partial_s u(s, t)$ and $\partial_t u(s, t)$. The Cauchy-Riemann equation (1.1) makes this easy, because

$$\partial_t u(s, t) = du(s, t)\partial_t = du(s, t)i\partial_s = i du(s, t)\partial_s = i \partial_s u(s, t),$$

which implies that $\partial_s u(s, t)$ and $\partial_t u(s, t)$ are orthogonal vectors of the same length. Thus the area of u is

$$\text{Area}(u) = \int_{\mathcal{U}} |\partial_s u| |\partial_t u| ds dt = \frac{1}{2} \int_{\mathcal{U}} (|\partial_s u|^2 + |\partial_t u|^2) ds dt,$$

where we’ve used the fact that $|\partial_s u| = |\partial_t u|$ to write things slightly more symmetrically. Notice that the right hand side is really an *analytical* quantity: up to a constant it is the square of the L^2 -norm of the first derivative of u .

Let us now write this area in a slightly different, more topological way. If $\langle \cdot, \cdot \rangle$ denotes the standard Hermitian inner product on \mathbb{C}^n , notice that one can define a differential 2-form on \mathbb{R}^{2n} by the expression

$$\omega_0(X, Y) = \text{Re}\langle iX, Y \rangle.$$

Writing points in \mathbb{C}^n via the coordinates $(p_1 + iq_1, \dots, p_n + iq_n)$, one can show that ω_0 in these coordinates takes the form

$$(1.2) \quad \omega_0 = \sum_{j=1}^n dp_j \wedge dq_j.$$

Exercise 1.2. Prove (1.2), and then show that ω_0 has the following three properties:

- (1) It is *nondegenerate*: $\omega_0(V, \cdot) = 0$ for some vector V if and only if $V = 0$. Equivalently, for each $z \in \mathbb{R}^{2n}$, the map $T_z \mathbb{R}^{2n} \rightarrow T_z^* \mathbb{R}^{2n} : V \mapsto \omega_0(V, \cdot)$ is an isomorphism.
- (2) It is *closed*: $d\omega_0 = 0$.
- (3) The n -fold product $\omega_0^n = \omega_0 \wedge \dots \wedge \omega_0$ is a constant multiple of the natural volume form on \mathbb{R}^{2n} .

Exercise 1.3. Show that a 2-form ω on \mathbb{R}^{2n} (and hence on any $2n$ -dimensional manifold) is nondegenerate if and only if ω^n is a volume form.

Using ω_0 , we see that the area of the parallelogram above is also

$$|\partial_s u| \cdot |\partial_t u| = |\partial_t u|^2 = \text{Re}\langle \partial_t u, \partial_t u \rangle = \text{Re}\langle i\partial_s u, \partial_t u \rangle = \omega_0(\partial_s u, \partial_t u),$$

thus

$$(1.3) \quad \text{Area}(u) = \|du\|_{L^2}^2 = \int_{\mathcal{U}} u^* \omega_0.$$

This is the first appearance of symplectic geometry in our study of holomorphic curves; we call ω_0 the **standard symplectic form** on \mathbb{R}^{2n} . The point is that the expression on the right hand side of (1.3) is essentially topological: it depends only on the evaluation of a certain closed 2-form on the 2-chain defined by $u(\mathcal{U})$. The present example is trivial because we're only working in \mathbb{R}^{2n} , but as we'll see later in more interesting examples, one can often find an easy topological bound on this integral, which by (1.3) implies a bound on the analytical quantity $\|du\|_{L^2}^2$. One can use this to derive compactness results for spaces of holomorphic curves, which then encode symplectic topological information about the space in which these curves live. We'll come back to this theme again and again.

1.2. Hamiltonian systems and symplectic manifolds

To motivate the study of symplectic manifolds in general, let us see how symplectic structures arise naturally in classical mechanics. We shall only sketch the main ideas here; a good comprehensive introduction may be found in [Arn89].

Consider a mechanical system with “ n degrees of freedom” moving under the influence of a Newtonian potential V . This means there are n “position” variables $q = (q_1, \dots, q_n) \in \mathbb{R}^n$, which are functions of time t that satisfy the second order differential equation

$$(1.4) \quad m_i \ddot{q}_i = -\frac{\partial V}{\partial q_i},$$

where $m_i > 0$ are constants representing the masses of the various particles, and $V : \mathbb{R}^n \rightarrow \mathbb{R}$ is a smooth function, the “potential”. The space \mathbb{R}^n , through which the vector $q(t)$ moves, is called the **configuration space** of the system. The basic idea of Hamiltonian mechanics is to turn this 2nd order system into a 1st order system by introducing an extra set of “momentum” variables $p = (p_1, \dots, p_n) \in \mathbb{R}^n$, where $p_i = m_i \dot{q}_i$. The space \mathbb{R}^{2n} with coordinates (p, q) is then called **phase space**, and we define a real valued function on phase space called the **Hamiltonian**, by

$$H : \mathbb{R}^{2n} \rightarrow \mathbb{R} : (p, q) \mapsto \frac{1}{2} \sum_{i=1}^n \frac{p_i^2}{m_i} + V(q).$$

Physicists will recognize this as the “total energy” of the system, but its main significance in the present context is that the combination of the second order system (1.4) with our definition of p is now equivalent to the $2n$ first order equations,

$$(1.5) \quad \dot{q}_i = \frac{\partial H}{\partial p_i}, \quad \dot{p}_i = -\frac{\partial H}{\partial q_i}.$$

These are **Hamilton's equations** for motion in phase space.

The motion of $x(t) := (p(t), q(t))$ in \mathbb{R}^{2n} can be described in more geometric terms: it is an orbit of the vector field

$$(1.6) \quad X_H(p, q) = \sum_{i=1}^n \left(\frac{\partial H}{\partial p_i} \frac{\partial}{\partial q_i} - \frac{\partial H}{\partial q_i} \frac{\partial}{\partial p_i} \right).$$

As we'll see in a moment, vector fields of this form have some important properties that have nothing to do with our particular choice of the function H , thus it is

sensible to call any vector field defined by this formula (for an arbitrary smooth function $H : \mathbb{R}^{2n} \rightarrow \mathbb{R}$) a **Hamiltonian vector field**. This is where the symplectic structure enters the story.

Exercise 1.4. Show that the vector field X_H of (1.6) can be characterized as the unique vector field on \mathbb{R}^{2n} that satisfies $\omega_0(X_H, \cdot) = -dH$.

The above exercise shows that the symplectic structure makes it possible to write down a much simplified definition of the Hamiltonian vector field. Now we can already prove something slightly impressive.

Proposition 1.5. *The flow φ_H^t of X_H satisfies $(\varphi_H^t)^*\omega_0 = \omega_0$ for all t .*

PROOF. Using Cartan's formula for the Lie derivative of a form, together with the characterization of X_H in Exercise 1.4 and the fact that ω_0 is closed, we compute $\mathcal{L}_{X_H}\omega_0 = d\iota_{X_H}\omega_0 + \iota_{X_H}d\omega_0 = -d^2H = 0$. \square

By Exercise 1.2, one can compute volumes on \mathbb{R}^{2n} by integrating the n -fold product $\omega_0 \wedge \dots \wedge \omega_0$, thus an immediate consequence of Prop. 1.5 is the following:

Corollary 1.6 (Liouville's theorem). *The flow of X_H is volume preserving.*

Notice that in most of this discussion we've not used our precise knowledge of the 2-form ω_0 or function H . Rather, we've used the fact that ω_0 is nondegenerate (to characterize X_H via ω_0 in Exercise 1.4), and the fact that it's closed (in the proof of Prop. 1.5). It is therefore natural to generalize as follows.

Definitions 1.7. A **symplectic form** on a $2n$ -dimensional manifold M is a smooth differential 2-form ω that is both closed and nondegenerate. The pair (M, ω) is then called a **symplectic manifold**. Given a smooth function $H : M \rightarrow \mathbb{R}$, the corresponding **Hamiltonian vector field** is defined to be the unique vector field $X_H \in \text{Vec}(M)$ such that¹

$$(1.7) \quad \omega(X_H, \cdot) = -dH.$$

For two symplectic manifolds (M_1, ω_1) and (M_2, ω_2) , a smooth map $\varphi : M_1 \rightarrow M_2$ is called **symplectic** if $\varphi^*\omega_2 = \omega_1$. If φ is a symplectic embedding, then we say that $\varphi(M_1)$ is a **symplectic submanifold** of (M_2, ω_2) . If φ is symplectic and is also a diffeomorphism, it is called a **symplectomorphism**, and we then say that (M_1, ω_1) and (M_2, ω_2) are **symplectomorphic**.

Repeating verbatim the argument of Prop. 1.5, we see now that any Hamiltonian vector field on a symplectic manifold (M, ω) defines a smooth 1-parameter family of symplectomorphisms. If we define volumes on M by integrating the $2n$ -form ω^n (see Exercise 1.3), then all symplectomorphisms are volume preserving—in particular this applies to the flow of X_H .

Remark 1.8. An odd-dimensional manifold can never admit a nondegenerate 2-form. (Why not?)

¹Some sources in the literature define X_H by $\omega(X_H, \cdot) = dH$, in which case one must choose different sign conventions for the orientation of phase space and definition of ω_0 . One must always be careful not to mix sign conventions from different sources—that way you could prove anything!

1.3. Some favorite examples

We now give a few examples of symplectic manifolds (other than $(\mathbb{R}^{2n}, \omega_0)$) which will be useful to have in mind.

Example 1.9. Suppose N is any smooth n -manifold and (q_1, \dots, q_n) are a choice of coordinates on an open subset $\mathcal{U} \subset N$. These naturally define coordinates $(p_1, \dots, p_n, q_1, \dots, q_n)$ on the cotangent bundle $T^*\mathcal{U} \subset T^*N$, where an arbitrary cotangent vector at $q \in \mathcal{U}$ is expressed as

$$p_1 dq_1 + \dots + p_n dq_n.$$

Interpreted differently, this expression also defines a smooth 1-form on $T^*\mathcal{U}$; we abbreviate it by $p dq$.

Exercise 1.10. Show that the 1-form $p dq$ doesn't actually depend on the choice of coordinates (q_1, \dots, q_n) .

What the above exercise reveals is that T^*N globally admits a *canonical* 1-form λ , whose expression in the local coordinates (p, q) always looks like $p dq$. Moreover, $d\lambda$ is clearly a symplectic form, as it looks exactly like (1.2) in coordinates. We call this the canonical symplectic form on T^*N . Using this symplectic structure, the cotangent bundle can be thought of as the “phase space” of a smooth manifold, and is a natural setting for studying Hamiltonian systems when the configuration space is something other than a Euclidean vector space (e.g. a “constrained” mechanical system).

Example 1.11. On any oriented surface Σ , a 2-form ω is symplectic if and only if it is an area form, and the symplectomorphisms are precisely the area preserving diffeomorphisms. Observe that one can always find area preserving diffeomorphisms between small open subsets of (\mathbb{R}^2, ω_0) and (Σ, ω) , thus every point in Σ has a neighborhood admitting local coordinates (p, q) in which $\omega = dp \wedge dq$.

Example 1.12. A more interesting example of a closed symplectic manifold is the n -dimensional complex projective space $\mathbb{C}P^n$. This is both a real $2n$ -dimensional symplectic manifold and a complex n -dimensional manifold, as we will now show. By definition, $\mathbb{C}P^n$ is the space of complex lines in \mathbb{C}^{n+1} , which we can express in two equivalent ways as follows:

$$\mathbb{C}P^n = (\mathbb{C}^{n+1} \setminus \{0\})/\mathbb{C}^* = S^{2n+1}/S^1.$$

In the first case, we divide out the natural free action (by scalar multiplication) of the multiplicative group $\mathbb{C}^* := \mathbb{C} \setminus \{0\}$ on $\mathbb{C}^{n+1} \setminus \{0\}$, and the second case is the same thing but restricting to the unit sphere $S^{2n+1} \subset \mathbb{C}^{n+1} = \mathbb{R}^{2n+2}$ and unit circle $S^1 \subset \mathbb{C} = \mathbb{R}^2$. To define a symplectic form, consider first the 1-form λ on S^{2n+1} defined for $z \in S^{2n+1} \subset \mathbb{C}^{n+1}$ and $X \in T_z S^{2n+1} \subset \mathbb{C}^{n+1}$ by

$$\lambda_z(X) = \langle iz, X \rangle,$$

where $\langle \cdot, \cdot \rangle$ is the standard Hermitian inner product on \mathbb{C}^{n+1} . (Take a moment to convince yourself that this expression is always real.) Since λ is clearly invariant under the S^1 -action on S^{2n+1} , the same is true for the closed 2-form $d\lambda$, which therefore descends to a closed 2-form ω_0 on $\mathbb{C}P^n$.

Exercise 1.13. Show that ω_0 as defined above is symplectic.

The complex manifold structure of $\mathbb{C}P^n$ can be seen explicitly by thinking of points in $\mathbb{C}P^n$ as equivalence classes of vectors $(z_0, \dots, z_n) \in \mathbb{C}^{n+1} \setminus \{0\}$, with two vectors equivalent if they are complex multiples of each other. We will always write the equivalence class represented by $(z_0, \dots, z_n) \in \mathbb{C}^{n+1} \setminus \{0\}$ as

$$[z_0 : \dots : z_n] \in \mathbb{C}P^n.$$

Then for each $k = 0, \dots, n$, there is an embedding

$$(1.8) \quad \iota_k : \mathbb{C}^n \hookrightarrow \mathbb{C}P^n : (z_1, \dots, z_n) \mapsto [z_1 : \dots : z_{k-1} : 1 : z_k : \dots : z_n],$$

whose image is the complement of the subset

$$\mathbb{C}P^{n-1} \cong \{[z_1 : \dots : z_{k-1} : 0 : z_k : \dots : z_n] \in \mathbb{C}P^n \mid (z_1, \dots, z_n) \in \mathbb{C}^n\}.$$

Exercise 1.14. Show that if the maps ι_k^{-1} are thought of as complex coordinate charts on open subsets of $\mathbb{C}P^n$, then the transition maps $\iota_k^{-1} \circ \iota_j$ are all holomorphic.

By the exercise, $\mathbb{C}P^n$ naturally carries the structure of a complex manifold such that the embeddings $\iota_k : \mathbb{C}^n \rightarrow \mathbb{C}P^n$ are holomorphic. Each of these embeddings also defines a decomposition of $\mathbb{C}P^n$ into $\mathbb{C}^n \cup \mathbb{C}P^{n-1}$, where $\mathbb{C}P^{n-1}$ is a complex submanifold of (complex) codimension one. The case $n = 1$ is particularly enlightening, as here the decomposition becomes $\mathbb{C}P^1 = \mathbb{C} \cup \{\text{point}\} \cong S^2$; this is simply the Riemann sphere with its natural complex structure, where the “point at infinity” is $\mathbb{C}P^0$. In the case $n = 2$, we have $\mathbb{C}P^2 \cong \mathbb{C}^2 \cup \mathbb{C}P^1$, and we’ll occasionally refer to the complex submanifold $\mathbb{C}P^1 \subset \mathbb{C}P^2$ as the “sphere at infinity”.

We continue for a moment with the example of $\mathbb{C}P^n$ in order to observe that it contains an abundance of holomorphic spheres. Take for instance the case $n = 2$: then for any $\zeta \in \mathbb{C}$, we claim that the holomorphic embedding

$$u_\zeta : \mathbb{C} \rightarrow \mathbb{C}^2 : z \mapsto (z, \zeta)$$

extends naturally to a holomorphic embedding of $\mathbb{C}P^1$ in $\mathbb{C}P^2$. Indeed, using ι_2 to include \mathbb{C}^2 in $\mathbb{C}P^2$, $u_\zeta(z)$ becomes the point $[z : \zeta : 1] = [1 : \zeta/z : 1/z]$, and as $z \rightarrow \infty$, this converges to the point $x_0 := [1 : 0 : 0]$ in the sphere at infinity. One can check using alternate charts that this extension is indeed a holomorphic map. The collection of all these embeddings $u_\zeta : \mathbb{C}P^1 \rightarrow \mathbb{C}P^2$ thus gives a very nice decomposition of $\mathbb{C}P^2$: together with the sphere at infinity, they foliate the region $\mathbb{C}P^2 \setminus \{x_0\}$, but all intersect precisely at x_0 (see Figure 1). This decomposition will turn out to be crucial in the proof of Theorem 1.25, stated below.

1.4. Darboux’s theorem and the Moser deformation trick

In Riemannian geometry, two Riemannian manifolds of the same dimension with different metrics can have quite different local structures: there can be no isometries between them, not even *locally*, unless they have the same curvature. The following basic result of symplectic geometry shows that in the symplectic world, things are quite different. We will give a proof using the beautiful *Moser deformation*




FIGURE 1. $\mathbb{C}P^2 \setminus \{x_0\}$ is foliated by holomorphic spheres that all intersect at x_0 .

trick, which has several important applications throughout symplectic and contact geometry, as we'll soon see.²

Theorem 1.15 (Darboux's theorem). *Near every point in a symplectic manifold (M, ω) , there are local coordinates $(p_1, \dots, p_n, q_1, \dots, q_n)$ in which $\omega = \sum_i dp_i \wedge dq_i$.*

PROOF. Denote by $(p_1, \dots, p_n, q_1, \dots, q_n)$ the standard coordinates on \mathbb{R}^{2n} and define the standard symplectic form ω_0 by (1.2); this is the exterior derivative of the 1-form

$$\lambda_0 = \sum_j p_j dq_j.$$

Since the statement in the theorem is purely local, we can assume (by choosing local coordinates) that M is an open neighborhood of the origin in \mathbb{R}^{2n} , on which ω is any closed, nondegenerate 2-form. Then it will suffice to find two open neighborhoods $\mathcal{U}, \mathcal{U}_0 \subset \mathbb{R}^{2n}$ of 0, and a diffeomorphism

$$\varphi : \mathcal{U}_0 \rightarrow \mathcal{U}$$

preserving 0 such that $\varphi^*\omega = \omega_0$. Using Exercise 1.16 below (the “linear Darboux's theorem”), we can also assume after a linear change of coordinates that $\varphi^*\omega$ and ω_0 match at the origin.

The idea behind the Moser trick is now the following bit of optimism: we assume that the desired diffeomorphism φ is the time 1 flow of a time-dependent vector field defined near 0, and derive conditions that this vector field must satisfy. In fact, we will be a bit more ambitious: consider the smooth 1-parameter family of 2-forms

$$\omega_t = t\omega + (1-t)\omega_0, \quad t \in [0, 1]$$

which interpolate between ω_0 and ω . These are all obviously closed, and if we restrict to a sufficiently small neighborhood of the origin then they are near ω_0 and

²An alternative approach to Darboux's theorem may be found in [Arn89].

thus nondegenerate. Our goal is to find a time-dependent vector field Y_t on some neighborhood of 0, for $t \in [0, 1]$, whose flow φ_t is well defined on some smaller neighborhood of 0 and satisfies

$$\varphi_t^* \omega_t = \omega_0$$

for all $t \in [0, 1]$. Differentiating this expression with respect to t and writing $\dot{\omega}_t := \frac{\partial}{\partial t} \omega_t$, we find

$$\varphi_t^* \mathcal{L}_{Y_t} \omega_t + \varphi_t^* \dot{\omega}_t = 0,$$

which by Cartan's formula and the fact that ω_t is closed and φ_t is a diffeomorphism, implies

$$(1.9) \quad d\iota_{Y_t} \omega_t + \dot{\omega}_t = 0.$$

At this point it's useful to observe that if we restrict to a contractible neighborhood of the origin, ω (and hence also ω_t) is exact: let us write

$$\omega = d\lambda.$$

Moreover, by adding a constant 1-form, we can choose λ so that it matches λ_0 at the origin. Now if $\lambda_t := t\lambda + (1-t)\lambda_0$, we have $d\lambda_t = \omega_t$, and $\dot{\lambda}_t := \frac{\partial}{\partial t} \lambda_t = \lambda - \lambda_0$ vanishes at the origin. Plugging this into (1.9), we see now that it suffices to find a vector field Y_t satisfying

$$(1.10) \quad \omega_t(Y_t, \cdot) = -\dot{\lambda}_t.$$

Since ω_t is nondegenerate, this equation can be solved and determines a unique vector field Y_t , which vanishes at the origin since $\dot{\lambda}_t$ does. The flow φ_t therefore exists for all $t \in [0, 1]$ on a sufficiently small neighborhood of the origin, and φ_1 is the desired diffeomorphism. \square

Exercise 1.16. The following linear version of Darboux's theorem is an easy exercise in linear algebra and was the first step in the proof above: show that if Ω is any nondegenerate, antisymmetric bilinear form on \mathbb{R}^{2n} , then there exists a basis $(X_1, \dots, X_n, Y_1, \dots, Y_n)$ such that

$$\Omega(X_i, Y_i) = 1$$

and Ω vanishes on all other pairs of basis vectors. This is equivalent to the statement that \mathbb{R}^{2n} admits a linear change of coordinates in which Ω looks like the standard symplectic form ω_0 .

It's worth pointing out the crucial role played in the above proof by the relation (1.10), which is almost the same as the relation used to define Hamiltonian vector fields (1.7). The latter, together with the argument of Prop. 1.5, tells us that the group of symplectomorphisms on a symplectic manifold is fantastically large, as it contains all the flows of Hamiltonian vector fields, which are determined by arbitrary smooth real valued functions. For much the same reason, one can also always find an abundance of symplectic local coordinate charts (usually called *Darboux coordinates*). Contrast this with the situation on a Riemannian manifold, where the group of isometries is generally finite dimensional, and different metrics are usually not locally equivalent, but are distinguished by their curvature.

In light of Darboux's theorem, we can now give the following equivalent definition of a symplectic manifold:

Definition 1.17. A symplectic manifold is a $2n$ -dimensional manifold M together with an atlas of coordinate charts whose transition maps are symplectic (with respect to the standard symplectic structure of \mathbb{R}^{2n}).

In physicists' language, a symplectic manifold is thus a manifold that can be identified locally with Hamiltonian phase space, in the sense that all coordinate changes leave the form of Hamilton's equations unaltered.

Let us state one more important application of the Moser trick, this time of a more global nature. Recall that two symplectic manifolds (M, ω) and (M', ω') are called *symplectomorphic* if there exists a symplectomorphism between them, i.e. a diffeomorphism $\varphi : M \rightarrow M'$ such that $\varphi^*\omega' = \omega$. Working on a single manifold M , we say similarly that two symplectic structures ω and ω' are symplectomorphic³ if (M, ω) and (M, ω') are symplectomorphic. This is the most obvious notion of equivalence for symplectic structures, but there are others that are also worth considering.

Definition 1.18. Two symplectic structures ω and ω' on M are called **isotopic** if there is a symplectomorphism $(M, \omega) \rightarrow (M, \omega')$ that is isotopic to the identity.

Definition 1.19. Two symplectic structures ω and ω' on M are called **deformation equivalent** if M admits a **symplectic deformation** between them, i.e. a smooth family of symplectic forms $\{\omega_t\}_{t \in [0,1]}$ such that $\omega_0 = \omega$ and $\omega_1 = \omega'$. Similarly, two symplectic manifolds (M, ω) and (M', ω') are deformation equivalent if there exists a diffeomorphism $\varphi : M \rightarrow M'$ such that ω and $\varphi^*\omega'$ are deformation equivalent.

It is clear that if two symplectic forms are isotopic then they are also both symplectomorphic and deformation equivalent. It is not true, however, that a symplectic deformation always gives rise to an isotopy: one should not expect this, as isotopic symplectic forms on M must always represent the same cohomology class in $H_{\text{dR}}^2(M)$, whereas the cohomology class can obviously vary under general deformations. The remarkable fact is that this necessary condition is also sufficient!

Theorem 1.20 (Moser's stability theorem). *Suppose M is a closed manifold with a smooth 1-parameter family of symplectic forms $\{\omega_t\}_{t \in [0,1]}$ which all represent the same cohomology class in $H_{\text{dR}}^2(M)$. Then there exists a smooth isotopy $\{\varphi_t : M \rightarrow M\}_{t \in [0,1]}$, with $\varphi_0 = \text{Id}$ and $\varphi_t^*\omega_t = \omega_0$.*

Exercise 1.21. Use the Moser isotopy trick to prove the theorem. *Hint: in the proof of Darboux's theorem, we had to use the fact that symplectic forms are locally exact in order to get from (1.9) to (1.10). Here you will find the cohomological hypothesis helpful for the same reason. If you get stuck, see [MS98].*

Exercise 1.22. Show that if ω and ω' are two deformation equivalent symplectic forms on $\mathbb{C}P^n$, then ω is isotopic to $c\omega'$ for some constant $c > 0$.

³The words "isomorphic" and "diffeomorphic" can also be used here as synonyms.

1.5. From symplectic geometry to symplectic topology

As a consequence of Darboux’s theorem, symplectic manifolds have no local invariants—there is no “local symplectic geometry”. Globally things are different, and here there are a number of interesting questions one can ask, all of which fall under the heading of *symplectic topology*. (The word “topology” is used to indicate the importance of global rather than local phenomena.)

The most basic such question concerns the classification of symplectic structures. One can ask, for example, whether there exists a symplectic manifold (M, ω) that is diffeomorphic to \mathbb{R}^4 but not symplectomorphic to (\mathbb{R}^4, ω_0) , i.e. an “exotic” symplectic \mathbb{R}^4 . The answer turns out to be yes—exotic \mathbb{R}^{2n} ’s exist in fact for all n , see [ALP94]—but it changes if we prescribe the behavior of ω at infinity. The following result says that $(\mathbb{R}^{2n}, \omega_0)$ is actually the only aspherical symplectic manifold that is “standard at infinity”.

Theorem 1.23 (Gromov [Gro85]). *Suppose (M, ω) is a symplectic 4-manifold with $\pi_2(M) = 0$, and there are compact subsets $K \subset M$ and $\Omega \subset \mathbb{R}^4$ such that $(M \setminus K, \omega)$ and $(\mathbb{R}^4 \setminus \Omega, \omega_0)$ are symplectomorphic. Then (M, ω) is symplectomorphic to (\mathbb{R}^4, ω_0) .*

In a later chapter we will be able to prove a stronger version of this statement, as a corollary of some classification results for symplectic fillings of contact manifolds (cf. Theorem 1.57).

Another interesting question is the following: suppose (M_1, ω_1) and (M_2, ω_2) are symplectic manifolds of the same dimension $2n$, possibly with boundary, such that there exists a smooth embedding $M_1 \hookrightarrow M_2$. Can one also find a *symplectic* embedding $(M_1, \omega_1) \hookrightarrow (M_2, \omega_2)$? What phenomena related to the symplectic structures can prevent this? There’s one obstruction that jumps out immediately: there can be no such embedding unless

$$\int_{M_1} \omega_1^n \leq \int_{M_2} \omega_2^n,$$

i.e. M_1 has no more volume than M_2 . In dimension two there’s nothing more to say, because symplectic and area preserving maps are the same thing. But in dimension $2n$ for $n \geq 2$, it was not known for a long time whether there are obstructions to symplectic embeddings other than the volume. A good thought experiment along these lines is the “squeezing” question: denote by B_r^{2n} the ball of radius r about the origin in \mathbb{R}^{2n} . Then it’s fairly obvious that for any $r, R > 0$ one can always find a volume preserving embedding

$$B_r^{2n} \hookrightarrow B_R^2 \times \mathbb{R}^{2n-2},$$

even if $r > R$, for then one can “squeeze” the first two dimensions of B_r^{2n} into B_R^2 but make up for it by spreading out further in \mathbb{R}^{2n-2} . But can one do this *symplectically*? The answer was provided by the following groundbreaking result:

Theorem 1.24 (Gromov’s “nonsqueezing” theorem [Gro85]). *There exists a symplectic embedding of (B_r^{2n}, ω_0) into $(B_R^2 \times \mathbb{R}^{2n-2}, \omega_0)$ if and only if $r \leq R$.*

This theorem was one of the first important applications of pseudoholomorphic curves. We will prove it in Chapter 5, and will spend a great deal of time in the

next few chapters learning the technical machinery that is needed to understand the proof.

We will close this brief introduction to symplectic topology by sketching the proof of a result that was introduced in [Gro85] and later generalized by McDuff, and provides us with a good excuse to introduce J -holomorphic curves. Recall from §1.3 that $\mathbb{C}P^2$ admits a singular foliation by embedded spheres that all intersect each other at one point, and all can be parametrized by holomorphic maps $\mathbb{C}P^1 \rightarrow \mathbb{C}P^2$. One can check that these spheres are also symplectic submanifolds with respect to the standard symplectic structure ω_0 introduced in Example 1.12; moreover, they intersect each other positively, so their self-intersection numbers are always 1. The following result essentially says that the existence of such a symplectically embedded sphere is a rare phenomenon: it can only occur in a very specific set of symplectic 4-manifolds, of which $(\mathbb{C}P^2, \omega_0)$ is the simplest. It also illustrates an important feature of symplectic topology specifically in four dimensions: once you find a single holomorphic curve with sufficiently nice local properties, it can sometimes fully determine the manifold in which it lives.

Theorem 1.25 (M. Gromov [Gro85] and D. McDuff [McD90]). *Suppose (M, ω) is a closed and connected symplectic 4-manifold containing a symplectically embedded 2-sphere $C \subset M$ with self-intersection $C \bullet C = 1$, but no symplectically embedded 2-sphere with self-intersection -1 . Then (M, ω) is symplectomorphic to $(\mathbb{C}P^2, c\omega_0)$, where $c > 0$ is a constant and ω_0 is the standard symplectic form on $\mathbb{C}P^2$.*

The idea of the proof is to choose appropriate data so that the symplectic submanifold $C \subset M$ can be regarded in some sense as a holomorphic curve, and then analyze the global structure of the space of holomorphic curves to which it belongs. It turns out that for a combination of analytical and topological reasons, this space will contain a smooth family of embedded holomorphic spheres that fill all of M and all intersect each other at one point, thus reproducing the singular foliation of Figure 1. This type of decomposition is a well known object in algebraic geometry and has more recently become quite popular in symplectic topology as well: it's called a *Lefschetz pencil*. As we'll see when we generalize Theorem 1.25 in a later chapter, there is an intimate connection between isotopy classes of Lefschetz pencils and deformation classes of symplectic structures: in the present case, the existence of this Lefschetz pencil implies that (M, ω) is symplectically deformation equivalent to $(\mathbb{C}P^2, \omega_0)$, and thus also symplectomorphic due to the Moser stability theorem (see Exercise 1.22).

The truly nontrivial part of the proof is the analysis of the moduli space of holomorphic curves, and this is what we'll concentrate on for the next several chapters. As a point of departure, consider the formulation (1.1) of the Cauchy-Riemann equations at the beginning of this chapter. Here u was a map from an open subset of \mathbb{C}^m into \mathbb{C}^n , but one can also make sense of (1.1) when u is a map between two complex manifolds. In such a situation, u is called holomorphic if and only if it looks holomorphic in any choice of holomorphic local coordinates. To put this in coordinate-free language, the tangent spaces of any complex manifold X are naturally complex vector spaces, on which multiplication by i makes sense, thus defining

a natural bundle endomorphism

$$i : TX \rightarrow TX$$

that satisfies $i^2 = -\mathbb{1}$. Then (1.1) makes sense globally and is the equation defining holomorphic maps between any two complex manifolds.

In the present situation, we're interested in smooth maps $u : \mathbb{CP}^1 \rightarrow M$. The domain is thus a complex manifold, but the target might not be, which means we lack an ingredient needed to write down the right hand side of (1.1). It turns out that one can always find an object to fill this role, i.e. a fiberwise linear map $J : TM \rightarrow TM$ with the following properties:

- $J^2 = -\mathbb{1}$,
- $\omega(\cdot, J\cdot)$ defines a Riemannian metric on M .

The first condition allows us to interpret J as “multiplication by i ”, thus turning the tangent spaces of M into complex vector spaces. The second reproduces the relation between i and ω_0 that exists in \mathbb{R}^{2n} , thus generalizing the important interaction between *symplectic* and *complex* that we illustrated in §1.1: complex subspaces of TM are also symplectic, and their areas can be computed in terms of ω . These conditions make J into a *compatible almost complex structure* on (M, ω) ; we will show by fairly elementary methods in §2.2 that such objects always exist on symplectic manifolds, in fact there are many of them. Now, the fact that C is embedded in M *symplectically* also allows us to arrange the following additional condition:

- the tangent spaces $TC \subset TM$ are invariant under J .

We are thus ready to introduce the following generalization of the Cauchy-Riemann equation: consider smooth maps $u : \mathbb{CP}^1 \rightarrow M$ whose differential is a complex linear map at every point, i.e.

$$(1.11) \quad Tu \circ i = J \circ Tu.$$

Solutions to (1.11) are called *pseudoholomorphic*, or more specifically, *J-holomorphic* spheres in M . Now pick a point $x_0 \in C$ and consider the following space of *J-holomorphic* spheres,

$$\begin{aligned} \mathcal{M} &:= \{u \in C^\infty(\mathbb{CP}^1, M) \mid Tu \circ i = J \circ Tu, \\ &\quad u_*[\mathbb{CP}^1] = [C] \in H_2(M), \\ &\quad u(0) = x_0\} / \sim, \end{aligned}$$

where $u \sim u'$ if there is a holomorphic diffeomorphism $\varphi : \mathbb{CP}^1 \rightarrow \mathbb{CP}^1$ such that $u' = u \circ \varphi$ and $\varphi(0) = 0$. We assign to \mathcal{M} the natural topology defined by C^∞ -convergence of smooth maps $\mathbb{CP}^1 \rightarrow M$.

Lemma 1.26. *\mathcal{M} is not empty: in particular it contains an embedded *J-holomorphic* sphere whose image is C .*

PROOF. Since C has J -invariant tangent spaces, any diffeomorphism $u_0 : \mathbb{CP}^1 \rightarrow C$ with $u_0(0) = x_0$ allows us to pull back J to an almost complex structure $j := u_0^*J$ on \mathbb{CP}^1 . As we'll review in Chapter 4, the uniqueness of complex structures on S^2 then allows us to find a diffeomorphism $\varphi : \mathbb{CP}^1 \rightarrow \mathbb{CP}^1$ such that $\varphi(0) = 0$ and $\varphi^*j = i$, thus the desired curve is $u := u_0 \circ \varphi$. \square

The rest of the work is done by the following rather powerful lemma, which describes the global structure of \mathcal{M} . Its proof requires a substantial volume of analytical machinery which we will develop in the coming chapters; note that since M is not a complex manifold, the methods of complex analysis play only a minor role in this machinery, and are subsumed in particular by the theory of nonlinear elliptic PDEs. This is the point where we need the technical assumptions that $C \bullet C = 1$ and M contains no symplectic spheres of self-intersection -1 ,⁴ as such topological conditions figure into the index computations that determine the local structure of \mathcal{M} .

Lemma 1.27. *\mathcal{M} is compact and admits the structure of a smooth 2-dimensional manifold. Moreover, the curves in \mathcal{M} are all embeddings that do not intersect each other except at the point x_0 ; in particular, they foliate $M \setminus \{x_0\}$.*

By this result, the curves in \mathcal{M} form the fibers of a symplectic Lefschetz pencil on (M, ω) , so that the latter's diffeomorphism and symplectomorphism type are completely determined by the moduli space of holomorphic curves.

1.6. Contact geometry and the Weinstein conjecture

Contact geometry is often called the “odd-dimensional cousin” of symplectic geometry, and one context in which it arises naturally is in the study of Hamiltonian dynamics. Again we shall only sketch the main ideas; the book [HZ94] is recommended for a more detailed account.

Consider a $2n$ -dimensional symplectic manifold (M, ω) with a Hamiltonian $H : M \rightarrow \mathbb{R}$. By the definition of the Hamiltonian vector field, $dH(X_H) = -\omega(X_H, X_H) = 0$, thus the flow of X_H preserves the level sets

$$S_c := H^{-1}(c)$$


for $c \in \mathbb{R}$. If c is a regular value of H then S_c is a smooth manifold of dimension $2n - 1$, called a **regular energy surface**, and X_H restricts to a nowhere zero vector field on S_c .

Exercise 1.28. If $S_c \subset M$ is a regular energy surface, show that the direction of X_H is uniquely determined by the condition $\omega(X_H, \cdot)|_{TS_c} = 0$.

The directions in Exercise 1.28 define the so-called **characteristic line field** on S_c : its existence implies that the paths traced out on S_c by orbits of X_H depend only on S_c and on the symplectic structure, not on H itself. In particular, a closed orbit of X_H on S_c is merely a closed integral curve of the characteristic line field. It is thus meaningful to ask the following question:

Question. *Given a symplectic manifold (M, ω) and a smooth hypersurface $S \subset M$, does the characteristic line field on S have any closed integral curves?*

⁴As we'll see, the assumption of no symplectic spheres with self-intersection -1 is a surprisingly weak one: it can always be attained by modifying (M, ω) in a standard way known as “blowing down”.

FIGURE 2. A star-shaped hypersurface in \mathbb{R}^2 .

We shall often refer to closed integral curves of the characteristic line field on $S \subset M$ simply as *closed orbits on S* . There are examples of Hamiltonian systems that have no closed orbits at all, cf. [HZ94, §4.5]. However, the following result (and the related result of A. Weinstein [Wei78] for convex energy surfaces) singles out a special class of hypersurfaces for which the answer is always yes:

Theorem 1.29 (P. Rabinowitz [Rab78]). *Every star-shaped hypersurface in the standard symplectic \mathbb{R}^{2n} admits a closed orbit.*

Recall that a hypersurface $S \subset \mathbb{R}^{2n}$ is called **star-shaped** if it doesn't intersect the origin and the projection $\mathbb{R}^{2n} \setminus \{0\} \rightarrow S^{2n-1} : z \mapsto z/|z|$ restricts to a diffeomorphism $S \rightarrow S^{2n-1}$ (see Figure 2). In particular, S is then transverse to the radial vector field

$$(1.12) \quad Y := \frac{1}{2} \sum_{i=1}^n \left(p_i \frac{\partial}{\partial p_i} + q_i \frac{\partial}{\partial q_i} \right).$$

Exercise 1.30. Show that the vector field Y of (1.12) satisfies $\mathcal{L}_Y \omega_0 = \omega_0$.

Definition 1.31. A vector field Y on a symplectic manifold (M, ω) is called a **Liouville vector field** if it satisfies $\mathcal{L}_Y \omega = \omega$.

By Exercise 1.30, star-shaped hypersurfaces in \mathbb{R}^{2n} are always transverse to a Liouville vector field, and this turns out to be a very special property.

Definition 1.32. A hypersurface S in a symplectic manifold (M, ω) is said to be of **contact type** if some neighborhood of S admits a Liouville vector field that is transverse to S .

Given a closed contact type hypersurface $S \subset (M, \omega)$, one can use the flow of the Liouville vector field Y to produce a very nice local picture of (M, ω) near S . Define a 1-form on S by

$$\alpha = \iota_Y \omega|_S,$$

and choose $\epsilon > 0$ sufficiently small so that

$$\Phi : (-\epsilon, \epsilon) \times S \rightarrow M : (t, x) \mapsto \varphi_Y^t(x)$$

is an embedding, where φ_Y^t denotes the flow of Y .

Exercise 1.33.

- (1) Show that the flow of Y “dilates” the symplectic form, i.e. $(\varphi_Y^t)^*\omega = e^t\omega$.
- (2) Show that $\Phi^*\omega = d(e^t\alpha)$, where we define α as a 1-form on $(-\epsilon, \epsilon) \times S$ by pulling it back through the natural projection to S . *Hint: show first that if $\lambda := \iota_Y\omega$, then $\Phi^*\lambda = e^t\alpha$, and notice that $d\lambda = \omega$ by the definition of a Liouville vector field.*
- (3) Show that $d\alpha$ restricts to a nondegenerate skew-symmetric 2-form on the hyperplane field $\xi := \ker \alpha$ over S . As a consequence, ξ is transverse to a smooth line field ℓ on S characterized by the property that $X \in \ell$ if and only if $d\alpha(X, \cdot) = 0$.
- (4) Show that on each of the hypersurfaces $\{c\} \times S$ for $c \in (-\epsilon, \epsilon)$, the line field ℓ defined above is the characteristic line field with respect to the symplectic form $d(e^t\alpha)$.

Several interesting consequences follow from Exercise 1.33. In particular, the use of a Liouville vector field to identify a neighborhood of S with $(-\epsilon, \epsilon) \times S$ gives us a smooth family of hypersurfaces $S_c := \{c\} \times S$ whose characteristic line fields all have exactly the same dynamics. This provides some intuitive motivation to believe Theorem 1.29: it’s sufficient to find one hypersurface in the family S_c that admits a periodic orbit, for then they all do. As it turns out, one can prove a variety of “almost existence” results in 1-parameter families of hypersurfaces, e.g. in $(\mathbb{R}^{2n}, \omega_0)$, a result of Hofer-Zehnder [HZ90] and Struwe [Str90] implies that for any smooth 1-parameter family of hypersurfaces, almost every (in a measure theoretic sense) hypersurface in the family admits a closed orbit. This gives a proof of the following generalization of Theorem 1.29:

Theorem 1.34 (C. Viterbo [Vit87]). *Every contact type hypersurface in $(\mathbb{R}^{2n}, \omega_0)$ admits a closed orbit.*

Having generalized this far, it’s natural to wonder whether the crucial properties of a contact hypersurface can be considered independently of its embedding into a symplectic manifold. The answer comes from the 1-form α and hyperplane distribution $\xi = \ker \alpha \subset TS$ in Exercise 1.33.

Definition 1.35. A **contact form** on a $(2n - 1)$ -dimensional manifold is a smooth 1-form α such that $d\alpha$ is always nondegenerate on $\xi := \ker \alpha$. The hyperplane distribution ξ is then called a **contact structure**.

Exercise 1.36. Show that the condition of $d\alpha$ being nondegenerate on $\xi = \ker \alpha$ is equivalent to $\alpha \wedge (d\alpha)^{n-1}$ being a volume form on S , and that ξ is nowhere integrable if this is satisfied.

Given an orientation of S , we call the contact structure $\xi = \ker \alpha$ **positive** if the orientation induced by $\alpha \wedge (d\alpha)^{n-1}$ agrees with the given orientation. One can show that if $S \subset (M, \omega)$ is a contact type hypersurface with the natural orientation induced from M and a transverse Liouville vector field, then the induced contact structure is always positive.

Note that Liouville vector fields are far from unique, in fact:

Exercise 1.37. Show that if Y is a Liouville vector field on (M, ω) and X_H is any Hamiltonian vector field, then $Y + X_H$ is also a Liouville vector field.

Thus the contact form $\alpha = \iota_Y \omega|_S$ induced on a contact type hypersurface should not be considered an intrinsic property of the hypersurface. As the next result indicates, the contact *structure* is the more meaningful object.

Proposition 1.38. *Up to isotopy, the contact structure $\xi = \ker \alpha$ induced on a contact type hypersurface $S \subset (M, \omega)$ by $\alpha = \iota_Y \omega|_S$ is independent of the choice of Y .*

The proof of this is a fairly easy exercise using a standard fundamental result of contact geometry:

Theorem 1.39 (Gray's stability theorem). *If S is a closed $(2n - 1)$ -dimensional manifold and $\{\xi_t\}_{t \in [0,1]}$ is a smooth 1-parameter family of contact structures on S , then there exists a smooth 1-parameter family of diffeomorphisms $\{\varphi_t\}_{t \in [0,1]}$ such that $\varphi_0 = \text{Id}$ and $(\varphi_t)_* \xi_0 = \xi_t$.*

This is yet another application of the Moser deformation trick; we'll explain the proof at the end of this section. Note that the theorem provides an isotopy between any two deformation equivalent contact *structures*, but there is no such result for contact *forms*—that's one of the reasons why contact structures are considered to be more geometrically natural objects.

By now we hopefully have sufficient motivation to study odd-dimensional manifolds with contact structures. The pair (S, ξ) is called a **contact manifold**, and for two contact manifolds (S_1, ξ_1) and (S_2, ξ_2) of the same dimension, a smooth embedding $\varphi : S_1 \hookrightarrow S_2$ is called a **contact embedding**

$$(S_1, \xi_1) \hookrightarrow (S_2, \xi_2)$$

if $\varphi_* \xi_1 = \xi_2$. If φ is also a diffeomorphism, then we call it a **contactomorphism**. One of the main questions in contact topology is how to distinguish closed contact manifolds that aren't contactomorphic. We'll touch upon this subject in the next section.

But first there is more to say about Hamiltonian dynamics. We saw in Exercise 1.33 that the characteristic line field on a contact type hypersurface $S \subset (M, \omega)$ can be described in terms of a contact form α : it is the unique line field containing all vectors X such that $d\alpha(X, \cdot) = 0$, and is necessarily transverse to the contact structure. The latter implies that α is nonzero in this direction, so we can use it to choose a normalization, leading to the following definition.

Definition 1.40. Given a contact form α on a $(2n - 1)$ -dimensional manifold S , the **Reeb vector field** is the unique vector field X_α satisfying

$$d\alpha(X_\alpha, \cdot) = 0, \quad \text{and} \quad \alpha(X_\alpha) = 1.$$

Thus closed integral curves on contact hypersurfaces can be identified with closed orbits of their Reeb vector fields.⁵ The “intrinsic” version of Theorems 1.29 and 1.34 is then the following famous conjecture.

Conjecture 1.41 (Weinstein conjecture). *For every closed odd-dimensional manifold M with a contact form α , X_α has a closed orbit.*

The Weinstein conjecture is still open in general, though a proof in dimension three was produced recently by C. Taubes [Tau07], using Seiberg-Witten theory. Before this, there was a long history of partial results using the theory of pseudo-holomorphic curves, such as the following (see Definition 1.52 below for the definition of “overtwisted”):

Theorem 1.42 (Hofer [Hof93]). *Every Reeb vector field on a closed 3-dimensional overtwisted contact manifold admits a contractible periodic orbit.*

The key idea introduced in [Hof93] was to look at J -holomorphic curves for a suitable class of almost complex structures J in the so-called *symplectization* $(\mathbb{R} \times M, d(e^t \alpha))$ of a manifold M with contact form α . Since the symplectic form is now exact, it’s no longer useful to consider *closed* holomorphic curves, e.g. a minor generalization of (1.3) shows that all J -holomorphic spheres $u : \mathbb{C}P^1 \rightarrow \mathbb{R} \times M$ are constant:

$$\text{Area}(u) = \|du\|_{L^2}^2 = \int_{\mathbb{C}P^1} u^* d(e^t \alpha) = \int_{\partial \mathbb{C}P^1} u^*(e^t \alpha) = 0.$$

Instead, one considers J -holomorphic maps

$$u : \dot{\Sigma} \rightarrow \mathbb{R} \times M,$$

where $\dot{\Sigma}$ denotes a closed Riemann surface with finitely many *punctures*. It turns out that under suitable conditions, the image of u near each puncture approaches $\{\pm\infty\} \times M$ and becomes asymptotically close to a cylinder of the form $\mathbb{R} \times \gamma$, where γ is a closed orbit of X_α (see Figure 3). Thus an existence result for punctured holomorphic curves in $\mathbb{R} \times M$ implies the Weinstein conjecture on M .

To tie up a loose end, here’s the proof of Gray’s stability theorem, followed by another important contact application of the Moser trick.

⁵Note that since Liouville vector fields are not unique, the Reeb vector field on a contact hypersurface is not uniquely determined, but its *direction* is.




FIGURE 3. A three-punctured pseudoholomorphic torus in the symplectization of a contact manifold.

PROOF OF THEOREM 1.39. Assume S is a closed manifold with a smooth family of contact forms $\{\alpha_t\}_{t \in [0,1]}$ defining contact structures $\xi_t = \ker \alpha_t$. We want to find a time-dependent vector field Y_t whose flow φ_t satisfies

$$(1.13) \quad \varphi_t^* \alpha_t = f_t \alpha_0$$

for some (arbitrary) smooth 1-parameter family of functions $f_t : S \rightarrow \mathbb{R}$. Differentiating this expression and writing $\dot{f}_t := \frac{\partial}{\partial t} f_t$ and $\dot{\alpha}_t := \frac{\partial}{\partial t} \alpha_t$, we have

$$\varphi_t^* (\dot{\alpha}_t + \mathcal{L}_{Y_t} \alpha_t) = \dot{f}_t \alpha_0 = \frac{\dot{f}_t}{f_t} \varphi_t^* \alpha_t,$$

and thus

$$(1.14) \quad \dot{\alpha}_t + d\iota_{Y_t} \alpha_t + \iota_{Y_t} d\alpha_t = g_t \alpha_t,$$

where we define a new family of functions $g_t : S \rightarrow \mathbb{R}$ via the relation

$$(1.15) \quad g_t \circ \varphi_t = \frac{\dot{f}_t}{f_t} = \frac{\partial}{\partial t} \log f_t.$$

Now to make life a bit simpler, we assume (optimistically!) that Y_t is always tangent to ξ_t , hence $\alpha_t(Y_t) = 0$ and the second term in (1.14) vanishes. We therefore need to find a vector field Y_t and function g_t such that

$$(1.16) \quad d\alpha_t(Y_t, \cdot) = -\dot{\alpha}_t + g_t \alpha_t.$$

Plugging in the Reeb vector field X_{α_t} on both sides, we find

$$0 = -\dot{\alpha}_t(X_{\alpha_t}) + g_t,$$

which determines the function g_t . Now restricting both sides of (1.16) to ξ_t , there is a unique solution for Y_t since $d\alpha_t|_{\xi_t}$ is nondegenerate. We can then integrate this vector field to obtain a family of diffeomorphisms φ_t , and integrate (1.15) to obtain f_t so that (1.13) is satisfied. \square

Exercise 1.43. Try to adapt the above argument to construct an isotopy such that $\varphi_t^* \alpha_t = \alpha_0$ for any two deformation equivalent contact forms. But don't try very hard.

Finally, just as there is no local symplectic geometry, there is no local contact geometry either:

Theorem 1.44 (Darboux's theorem for contact manifolds). *Near every point in a $(2n + 1)$ -dimensional manifold S with contact form α , there are local coordinates $(p_1, \dots, p_n, q_1, \dots, q_n, z)$ in which $\alpha = dz + \sum_i p_i dq_i$.*

Exercise 1.45. Prove the theorem using a Moser argument. If you get stuck, see [Gei08].

1.7. Symplectic fillings of contact manifolds

In the previous section, contact manifolds were introduced as objects that occur naturally as hypersurfaces in symplectic manifolds. In particular, every contact manifold (M, ξ) with contact form α is obviously a contact type hypersurface in its own symplectization $(\mathbb{R} \times M, d(e^t \alpha))$, though this example is in some sense trivial. By contrast, it is far from obvious whether any given contact manifold can occur as a contact hypersurface in a *closed* symplectic manifold, or relatedly, if it is a “contact type boundary” of some compact symplectic manifold.

Definition 1.46. A compact symplectic manifold (W, ω) with boundary is said to have **convex boundary** if there exists a Liouville vector field in a neighborhood of ∂W that points transversely out of ∂W .

Definition 1.47. A **strong symplectic filling** (also called a **convex filling**) of a closed contact manifold (M, ξ) is a compact symplectic manifold (W, ω) with convex boundary, such that ∂W with the contact structure induced by a Liouville vector field is contactomorphic to (M, ξ) .

Since we're now considering symplectic manifolds that are not closed, it's also possible for ω to be exact. Observe that a primitive λ of ω always gives rise to a Liouville vector field, since the unique vector field Y defined by $\iota_Y \omega = \lambda$ then satisfies

$$\mathcal{L}_Y \omega = d\iota_Y \omega = d\lambda = \omega.$$

Definition 1.48. A strong filling (W, ω) of (M, ξ) is called an **exact filling** if $\omega = d\lambda$ for some 1-form λ such that the vector field Y defined by $\iota_Y \omega = \lambda$ points transversely out of ∂W .

Exercise 1.49. Show that if (W, ω) is a compact symplectic manifold with boundary, Y is a Liouville vector field defined near ∂W and $\lambda = \iota_Y \omega$, then Y is positively transverse to ∂W if and only if $\lambda|_{\partial W}$ is a positive contact form.

The exercise makes possible the following alternative formulations of the above definitions:

- (1) A compact symplectic manifold (W, ω) with boundary is a strong filling if ∂W admits a contact form that extends to a primitive of ω on a neighborhood of ∂W .
- (2) A strong filling is exact if the primitive mentioned above can be extended globally over W .

- (3) A strong filling is exact if it has a transverse outward pointing Liouville vector field near ∂W that can be extended globally over W .

By now you're surely wondering what a "weak" filling is. Observe that for any strong filling (W, ω) with Liouville vector field Y and induced contact structure $\xi = \ker \iota_Y \omega$ on the boundary, ω has a nondegenerate restriction to ξ (see Exercise 1.33). The latter condition can be expressed without mentioning a Liouville vector field, hence:

Definition 1.50. A **weak symplectic filling** of a closed contact manifold (M, ξ) is a compact symplectic manifold (W, ω) with boundary, such that there exists a diffeomorphism $\varphi : \partial W \rightarrow M$ and ω has a nondegenerate restriction to $\varphi^* \xi$.

Remark 1.51. One important definition that we are leaving out of the present discussion is that of a *Stein filling*: this is a certain type of complex manifold with contact boundary, which is also an exact symplectic filling. The results we'll prove in these lectures for strong and exact fillings apply to Stein fillings as well, but we will usually not make specific mention of this since the Stein condition itself has no impact on our general setup. Much more on Stein manifolds can be found in the monographs [OS04] and [CE].

A contact manifold is called exactly/strongly/weakly fillable if it admits an exact/strong/weak filling. Recall that in the smooth category, every 3-manifold is the boundary of some 4-manifold; by contrast, we will see that many contact 3-manifolds are not symplectically fillable.

The unit ball in (\mathbb{R}^4, ω_0) obviously has convex boundary: the contact structure induced on S^3 is called the **standard** contact structure ξ_0 . But there are other contact structures on S^3 not contactomorphic to ξ_0 , and one way to see this is to show that they are not fillable. Indeed, it is easy (via "Lutz twists", see [Gei08] or [Gei06]) to produce a contact structure on S^3 that is *overtwisted*. Note that the following is not the standard definition⁶ of this term, but is equivalent due to a deep result of Eliashberg [Eli89].

Definition 1.52. A contact 3-manifold (M, ξ) is **overtwisted** if it admits a contact embedding of $(S^1 \times \mathbb{D}, \xi_{\text{OT}})$, where $\mathbb{D} \subset \mathbb{R}^2$ is the closed unit disk and ξ_{OT} is a contact structure of the form

$$\xi_{\text{OT}} = \ker [f(\rho) d\theta + g(\rho) d\phi]$$

with $\theta \in S^1$, (ρ, ϕ) denoting polar coordinates on \mathbb{D} , and $(f, g) : [0, 1] \rightarrow \mathbb{R}^2 \setminus \{0\}$ a smooth path that begins at $(1, 0)$ and winds counterclockwise around the origin, making at least one half turn.

For visualization, a portion of the domain $(S^1 \times \mathbb{D}, \xi_{\text{OT}})$ is shown in Figure 4. One of the earliest applications of holomorphic curves in contact topology was the following nonfillability result.

Theorem 1.53 (M. Gromov [Gro85] and Ya. Eliashberg [Eli90]). *If (M, ξ) is closed and overtwisted, then it is not weakly fillable.*

⁶It is standard to call a contact 3-manifold (M, ξ) overtwisted if it contains an embedded *overtwisted disk*, which is a disk $\mathcal{D} \subset M$ such that $T(\partial \mathcal{D}) \subset \xi$ but $T\mathcal{D}|_{\partial \mathcal{D}} \neq \xi|_{\partial \mathcal{D}}$.




FIGURE 4. An overtwisted contact structure.

The Gromov-Eliashberg proof worked by assuming a weak filling (W, ω) of (M, ξ) exists, then constructing a family of J -holomorphic disks in W with boundaries on a totally real submanifold in M and showing that this family leads to a contradiction if (M, ξ) contains an overtwisted disk. We will later present a proof that is similar in spirit but uses slightly different techniques: instead of dealing with boundary conditions for holomorphic disks, we will adopt Hofer's methods and consider punctured holomorphic curves in a noncompact symplectic manifold obtained by gluing a cylindrical end to ∂W . The advantage of this approach is that it generalizes nicely to prove the following related result on Giroux torsion, which is much more recent. Previous proofs due to D. Gay and Ghiggini and Honda required the large machinery of gauge theory and Heegaard Floer homology respectively, but we will only use punctured holomorphic curves.

Theorem 1.54 (D. Gay [Gay06], P. Ghiggini and K. Honda [GH]). *Suppose (M, ξ) is a closed contact 3-manifold that admits a contact embedding of $(T^2 \times [0, 1], \xi_T)$, where ξ_T is the contact structure defined in coordinates $(\theta, \phi, r) \in S^1 \times S^1 \times [0, 1]$ by*

$$\xi_T = \ker [\cos(2\pi r) d\theta + \sin(2\pi r) d\phi].$$

Then (M, ξ) is not strongly fillable. Moreover if the embedded torus $T^2 \times \{0\}$ separates M , then (M, ξ) is also not weakly fillable.

A contact 3-manifold that admits a contact embedding of $(T^2 \times [0, 1], \xi_T)$ as defined above is said to have **Giroux torsion**.

Example 1.55. Using coordinates $(\theta, \phi, \eta) \in S^1 \times S^1 \times S^1 = T^3$, one can define for each $N \in \mathbb{N}$ a contact structure $\xi_N = \ker \alpha_N$, where

$$\alpha_N = \cos(2\pi N\eta) d\theta + \sin(2\pi N\eta) d\phi.$$

Choosing the natural flat metric on $T^2 = S^1 \times S^1$, it's easy to show that the unit circle bundle in T^*T^2 is a contact type hypersurface contactomorphic to (T^3, ξ_1) , thus this is strongly (and even exactly) fillable. Giroux [Gir94] and Eliashberg [Eli96] have shown that (T^3, ξ_N) is in fact weakly fillable for all N , but Theorem 1.54 implies that it is not strongly fillable for $N \geq 2$ (a result originally proved by Eliashberg [Eli96]). Unlike the case of S^3 , none of these contact structures are overtwisted—one can see this easily from Theorem 1.42 and the exercise below.

Exercise 1.56. Derive expressions for the Reeb vector fields X_{α_N} on T^3 and show that none of them admit any contractible periodic orbits.

Finally, we mention one case of a fillable contact manifold in which all the symplectic fillings can be described quite explicitly. Earlier we defined the standard contact structure ξ_0 on S^3 to be the one that is induced on the convex boundary of a round ball in (\mathbb{R}^4, ω_0) . By looking at isotopies of convex boundaries and using Gray's stability theorem, you should easily be able to convince yourself that every star-shaped hypersurface in (\mathbb{R}^4, ω_0) has an induced contact structure isotopic to ξ_0 . Thus the regions bounded by these hypersurfaces, the “star-shaped domains” in (\mathbb{R}^4, ω_0) , can all be regarded as convex fillings of (S^3, ξ_0) . Are there any others? Well...

Theorem 1.57 (Eliashberg [Eli90]). *Every exact filling of (S^3, ξ_0) is symplectomorphic to a star-shaped domain in (\mathbb{R}^4, ω_0) .*

In fact we will just as easily be able to classify all the weak fillings of (S^3, ξ_0) up to symplectic deformation equivalence. Again, our proof will differ from Eliashberg's in using punctured holomorphic curves asymptotic to Reeb orbits instead of compact curves with totally real boundary conditions. But in either case, the proof has much philosophically in common with the proof of Theorem 1.25 that we already sketched: one first finds a single holomorphic curve, in this case near the boundary of the filling, and then lets the moduli space of such curves “spread out” until it yields a geometric decomposition of the filling.

CHAPTER 2

Local properties

2.1. Almost complex manifolds and J -holomorphic curves

We now begin the study of J -holomorphic curves in earnest by defining the nonlinear Cauchy-Riemann equation in its most natural setting, and then examining the analytical properties of its solutions. This will be the focus for the next few chapters.

Given a $2n$ -dimensional real vector space, we define a **complex structure** on V to be any linear map $J : V \rightarrow V$ such that $J^2 = -\mathbb{1}$. It's easy to see that a complex structure always exists when $\dim V$ is even, as one can choose a basis to identify V with \mathbb{R}^{2n} and identify this in turn with \mathbb{C}^n , so that the natural “multiplication by i ” on \mathbb{C}^n becomes a linear map on V . In the chosen basis, this linear map is represented by the matrix

$$J_0 := \begin{pmatrix} 0 & -1 & & \\ 1 & 0 & & \\ & & \ddots & \\ & & & 0 & -1 \\ & & & 1 & 0 \end{pmatrix}.$$

We call this the **standard** complex structure on \mathbb{R}^{2n} , and will alternately denote it by J_0 or i , depending on the context. A complex structure J on V allows us to view V as a complex n -dimensional vector space, in that we identify the scalar multiplication by any complex number $a+ib \in \mathbb{C}$ with the linear map $a\mathbb{1}+bJ$. A real linear map on V is then also complex linear in this sense if and only if it commutes with J . Similarly, we call a real linear map $A : V \rightarrow V$ **complex antilinear** if it anticommutes with J , i.e. $AJ = -JA$. This is equivalent to the requirement that A preserves vector addition but satisfies $A(\lambda v) = \bar{\lambda}Av$ for all $v \in V$ and complex scalars $\lambda \in \mathbb{C}$.

Exercise 2.1.

- (1) Show that for every even-dimensional vector space V with complex structure J , there exists a basis in which J takes the form of the standard complex structure J_0 .
- (2) Show that if V is an odd-dimensional vector space, then there is no linear map $J : V \rightarrow V$ satisfying $J^2 = -\mathbb{1}$.
- (3) Show that all real linear maps on \mathbb{R}^{2n} that commute with J_0 have positive determinant.

Note that due to the above exercise, a complex structure J on a $2n$ -dimensional vector space V induces a natural *orientation* on V , namely by defining any basis of the form $(v_1, Jv_1, \dots, v_n, Jv_n)$ to be positively oriented. This is equivalent to the statement that every complex finite dimensional vector space has a natural orientation as a *real* vector space.

The above notions can easily be generalized from spaces to bundles: if M is a topological space and $E \rightarrow M$ is a real vector bundle of even rank, then a **complex structure** on $E \rightarrow M$ is a smooth family of complex structures on the fibers of E , i.e. a section $J \in \Gamma(\text{End}(E))$ of the bundle $\text{End}(E)$ of fiber-preserving linear maps $E \rightarrow E$, such that $J^2 = -\mathbf{1}$. If $E \rightarrow M$ is a smooth vector bundle, then we will usually assume that J is also either smooth or of some specific differentiability class. A complex structure gives $E \rightarrow M$ the structure of a complex vector bundle, due to the following variation on Exercise 2.1 above.

Exercise 2.2.

- (1) Show that whenever $E \rightarrow M$ is a real vector bundle of even rank with a complex structure J , every point $p \in M$ lies in a neighborhood on which E admits a trivialization such that J takes the form of the standard complex structure J_0 .
- (2) Show that for any two trivializations having the above property, the transition map relating them is fiberwise complex linear (using the natural identification $\mathbb{R}^{2n} = \mathbb{C}^n$).

For this reason, it is often convenient to denote complex vector bundles of rank n as pairs (E, J) , where E is a real bundle of rank $2n$ and J is a complex structure on E . Note that not every real vector bundle of even rank admits a complex structure: the above discussion shows that such bundles must always be orientable, and this condition is not even generally sufficient except for the case of rank two.

For a smooth $2n$ -dimensional manifold M , we refer to any complex structure J on the tangent bundle TM as an **almost complex structure** on M , and the pair (M, J) is then an **almost complex manifold**. The reason for the word “almost” will be explained in a moment.

Example 2.3. Suppose M is a complex manifold of complex dimension n , i.e. there exist local charts covering M that identify subsets of M with subsets of \mathbb{C}^n such that all transition maps are holomorphic. Any choice of holomorphic local coordinates on a subset $\mathcal{U} \subset M$ then identifies the tangent spaces $T_p\mathcal{U}$ with \mathbb{C}^n . If we use this identification to assign the standard complex structure i to each tangent space $T_p\mathcal{U}$, then the fact that transition maps are holomorphic implies that this assignment doesn’t depend on the choice of coordinates (prove this!). Thus M has a natural almost complex structure J that looks like the standard complex structure in any holomorphic coordinate chart.

An almost complex structure is called **integrable** if it arises in the above manner from a system of holomorphic coordinate charts; in this case we drop the word “almost” and simply call J a **complex structure** on M . By definition then, a real manifold M admits a complex structure (i.e. an integrable almost complex structure)

if and only if it also admits coordinate charts that make it into a complex manifold. In contrast to Exercise 2.2, which applies to trivializations on vector bundles, one cannot always find a coordinate chart that makes a given almost complex structure look standard on a neighborhood. The following standard (but hard) result of complex analysis characterizes integrable complex structures; we include it here for informational purposes, but will not make essential use of it in the following.

Theorem 2.4. *The almost complex structure J on M is integrable if and only if the tensor N_J vanishes identically, where N_J is defined on two vector fields X and Y by*

$$(2.1) \quad N_J(X, Y) = [JX, JY] - J[JX, Y] - J[X, JY] - [X, Y].$$

The tensor (2.1) is called the **Nijenhuis tensor**.

Exercise 2.5.

- (1) Verify that (2.1) defines a tensor.
- (2) Show that N_J always vanishes if $\dim M = 2$.
- (3) Prove one direction of Theorem 2.4: if J is integrable, then N_J vanishes.

The converse direction is much harder to prove, see for instance [DK90, Chapter 2]. But if you believe this, then Exercise 2.5 has the following nice consequence:

Theorem 2.6. *Every almost complex structure on a surface is integrable.*

In other words, complex 1-dimensional manifolds are the same thing as almost complex manifolds of real dimension two. This theorem follows from an existence result for local pseudoholomorphic curves which we'll prove in §2.10. Actually, that existence result can be thought of as the first step in the proof of Theorem 2.4. Complex manifolds in the lowest dimension have a special status, and deserve a special name:

Definition 2.7. A **Riemann surface** is a complex manifold of complex dimension one.

By Theorem 2.6, a Riemann surface can equivalently be regarded as a surface Σ with an almost complex structure j , and we will thus typically denote Riemann surfaces as pairs (Σ, j) .

Surfaces are the easy special case; in dimensions four and higher, (2.1) does not usually vanish, in fact it is *generically* nonzero, which shows that in some sense “generic” almost complex structures are not integrable. Thus in higher dimensions, integrable complex structures are very rigid objects—too rigid for our purposes, as it will turn out. For instance, there are real manifolds that do not admit complex structures but do admit almost complex structures. It will be most important for our purposes to observe that *symplectic* manifolds always admit almost complex structures that are “compatible” with the symplectic form in a certain geometric sense. We'll come back to this in §2.2 and make considerable use of it in later applications, but for now, we are only interested in the local properties of J -holomorphic curves, and for this it suffices to work in the more general context of almost complex manifolds.

Definition 2.8. Suppose (Σ, j) is a Riemann surface and (M, J) is an almost complex manifold. A smooth map $u : \Sigma \rightarrow M$ is called **J -holomorphic** (or **pseudo-holomorphic**) if its differential at every point is complex linear, i.e.

$$(2.2) \quad Tu \circ j = J \circ Tu.$$

Note that in general, the equation (2.2) makes sense if u is only of class C^1 (or more generally, of Sobolev class $W^{1,p}$) rather than smooth, but it will turn out to follow from elliptic regularity (see §2.5 and §2.10) that J -holomorphic curves are *always* smooth if J is smooth—we will therefore assume smoothness whenever convenient. Equation (2.2) is a nonlinear first-order PDE, often called the **nonlinear Cauchy-Riemann equation**. If you are not accustomed to PDEs expressed in geometric notation, you may prefer to view it as follows: choose holomorphic local coordinates $s + it$ on a subset of Σ , so $j\partial_s = \partial_t$ and $j\partial_t = -\partial_s$ (note that we're assuming the integrability of j). Then (2.2) is locally equivalent to the equation

$$(2.3) \quad \partial_s u + J(u)\partial_t u = 0.$$

Notation. We will sometimes write $u : (\Sigma, j) \rightarrow (M, J)$ to mean that $u : \Sigma \rightarrow M$ is a map satisfying (2.2). When the domain is the open unit ball $B \subset \mathbb{C}$ (or any other open subset of \mathbb{C}) and we say $u : B \rightarrow M$ is J -holomorphic without specifying the complex structure of the domain, then the standard complex structure is implied, i.e. u is a pseudoholomorphic map $(B, i) \rightarrow (M, J)$ and thus satisfies (2.3).

Note that the standard Cauchy-Riemann equation for maps $u : \mathbb{C} \rightarrow \mathbb{C}^n$ can be written as $\partial_s u + i\partial_t u = 0$, thus (2.3) can be viewed as a perturbation of this. In fact, due to Exercise 2.1 one can always choose coordinates near a point $p \in M$ so that $J(p)$ is identified with the standard complex structure; then in a sufficiently small neighborhood of p , (2.3) really is a *small* perturbation of the usual Cauchy-Riemann equation. We'll make considerable use of this perspective in the following. Here is a summary of the most important results we aim to prove in this chapter.

Theorem. Assume (M, J) is a smooth almost complex manifold. Then:

- (REGULARITY) Every map $u : \Sigma \rightarrow M$ of class C^1 solving the nonlinear Cauchy-Riemann equation (2.2) is smooth (cf. Theorem 2.60).
- (LOCAL EXISTENCE) For any $p \in M$ and $X \in T_p M$, there exists a neighborhood $\mathcal{U} \subset \mathbb{C}$ of the origin and a J -holomorphic map $u : \mathcal{U} \rightarrow M$ such that $u(0) = p$ and $\partial_s u(0) = X$ in standard coordinates $s + it \in \mathcal{U}$ (cf. Theorem 2.71).
- (CRITICAL POINTS) If $u : \Sigma \rightarrow M$ is a nonconstant J -holomorphic curve with a critical point $z \in \Sigma$, then there is a neighborhood $\mathcal{U} \subset \Sigma$ of z such that $u|_{\mathcal{U} \setminus \{z\}}$ is a k -to-1 immersion for some $k \in \mathbb{N}$ (cf. Corollary 2.33 and Theorem 2.88).
- (INTERSECTIONS) Suppose $u_1 : \Sigma_1 \rightarrow M$ and $u_2 : \Sigma_2 \rightarrow M$ are two nonconstant J -holomorphic curves with an intersection $u_1(z_1) = u_2(z_2)$. Then there exist neighborhoods $\mathcal{U}_1 \subset \Sigma_1$ and $\mathcal{U}_2 \subset \Sigma_2$ such that the images $u_1(\mathcal{U}_1 \setminus \{z_1\})$ and $u_2(\mathcal{U}_2 \setminus \{z_2\})$ are either identical or disjoint (cf. Theorem 2.87).

This theorem amounts to the statement that locally, J -holomorphic curves behave much the same way as holomorphic curves, i.e. the same as in the integrable case. But since J is usually not integrable, the methods of complex analysis cannot be applied here, and we will instead need to employ techniques from the theory of elliptic PDEs. As preparation, we'll derive the natural linearization of (2.2) and introduce the theory of linear Cauchy-Riemann operators, as well as some fundamental ideas of global analysis, all of which will be useful in the chapters to come.

2.2. Compatible and tamed almost complex structures

For any given even-dimensional manifold M , it is not always immediately clear whether an almost complex structure exists. If $\dim M = 2$ for instance, then this is true if and only if M is orientable, and in higher dimensions the question is more complicated. We will not address this question in full generality, but merely show in the present section that for the cases we are most interested in, namely for *symplectic* manifolds, the answer is exactly as we might hope. The results of this section are independent of the rest of the chapter, but they will become crucial once we discuss compactness results and applications, from Chapter 5 onwards.

Denote the space of smooth almost complex structures on a given smooth $2n$ -dimensional manifold M by

$$\mathcal{J}(M) = \{J \in \Gamma(\text{End}(TM)) \mid J^2 = -\mathbf{1}\},$$

and assign to this space the natural C^∞ -topology. In this section, we will assume (M, ω) is a $2n$ -dimensional symplectic manifold, and consider two special subspaces of $\mathcal{J}(M)$:

$$\mathcal{J}^\tau(M, \omega) := \{J \in \mathcal{J}(M) \mid \omega(X, JX) > 0 \text{ for all } X \neq 0\},$$

$$\mathcal{J}(M, \omega) := \{J \in \mathcal{J}(M) \mid g_J(X, Y) := \omega(X, JY) \text{ is a Riemannian metric}\}.$$

We say that J is **tamed by** ω if $J \in \mathcal{J}^\tau(M, \omega)$, and it is **compatible with** ω if $J \in \mathcal{J}(M, \omega)$; clearly $\mathcal{J}(M, \omega) \subset \mathcal{J}^\tau(M, \omega)$. The key property of tame almost complex structures is that every complex line in a tangent space is also a symplectic subspace, and hence every embedded J -holomorphic curve parametrizes a symplectic submanifold. The only reason this condition is weaker than ω -compatibility is that we don't require the bilinear form $(X, Y) \mapsto \omega(X, JY)$ to be symmetric, but one can still symmetrize it to define a Riemannian metric,

$$(2.4) \quad g_J(X, Y) := \frac{1}{2} [\omega(X, JY) + \omega(Y, JX)],$$

which is identical to the above definition in the case $J \in \mathcal{J}(M, \omega)$.

Exercise 2.9. Show that a tamed almost complex structure $J \in \mathcal{J}^\tau(M, \omega)$ is also ω -compatible if and only if ω is J -invariant, i.e. $\omega(JX, JY) = \omega(X, Y)$ for all X and Y .

Most simple examples of almost complex structures one can write down on symplectic manifolds are compatible: e.g. this is true for the standard structure $i = J_0$ on the standard $(\mathbb{R}^{2n}, \omega_0)$, and for any complex structure compatible with the canonical orientation on a symplectic surface. Since every almost complex structure looks

like the standard one at a point in appropriate coordinates, it is easy to see that every J is *locally* tamed by some symplectic structure: namely, if $J \in \mathcal{J}(B^{2n})$ with $J(0) = i$, then J is tamed by ω_0 on B_ϵ^{2n} for sufficiently small $\epsilon > 0$ since tameness is an open condition.

It turns out that for a given symplectic structure, compatible almost complex structures always exist, and the global topologies of $\mathcal{J}(M, \omega)$ and $\mathcal{J}^\tau(M, \omega)$ are conveniently simple. Since it will be useful later, let us state this fact a little bit more generally: suppose $(E, \omega) \rightarrow B$ is a smooth **symplectic vector bundle** of finite rank, i.e. a vector bundle whose fibers are equipped with symplectic structures that vary smoothly. It is straightforward to show that such a bundle admits local trivializations that identify every fiber symplectically with $(\mathbb{R}^{2n}, \omega_0)$; see [MS98]. Recall that a complex structure on E is a smooth bundle endomorphism $J \in \Gamma(\text{End}(E))$ such that $J^2 = -\mathbb{1}$, and one can then adapt the definition above to speak of complex structures on E that are tamed by or compatible with ω . Denote the resulting spaces of complex structures by $\mathcal{J}^\tau(E, \omega)$ and $\mathcal{J}(E, \omega)$ respectively.

Proposition 2.10. *For any finite rank symplectic vector bundle $(E, \omega) \rightarrow B$, the spaces $\mathcal{J}(E, \omega)$ and $\mathcal{J}^\tau(E, \omega)$ are both nonempty and contractible.*

PROOF. We'll prove the statement only for $\mathcal{J}(E, \omega)$, which is all we need for applications. The fact that $\mathcal{J}^\tau(E, \omega)$ is not empty follows immediately, and two proofs of the contractibility of this space may be found in [MS98, Prop. 2.51].

For compatible complex structures, the basic idea is as follows: let $\mathfrak{M}(E)$ denote the space of smooth bundle metrics on E . There is then a natural continuous map

$$\mathcal{J}(E, \omega) \rightarrow \mathfrak{M}(E) : J \mapsto g_J,$$

where $g_J := \omega(\cdot, J\cdot)$. We shall construct a continuous left inverse to this map, i.e. a continuous map

$$\Phi : \mathfrak{M}(E) \rightarrow \mathcal{J}(E, \omega)$$

such that $\Phi(g_J) = J$ for every $J \in \mathcal{J}(E, \omega)$. Then since $\mathfrak{M}(E)$ is a nonempty convex subset of a vector space and hence contractible, the identity map $J \mapsto \Phi(g_J)$ can be contracted to a point by contracting $\mathfrak{M}(E)$.

To construct the map Φ , observe that if $g \in \mathfrak{M}(E)$ happens to be of the form g_J for some $J \in \mathcal{J}(E, \omega)$, then it is related to J by $\omega \equiv g(J\cdot, \cdot)$. For more general metrics g , this relation still determines J as a linear bundle map on E , and the latter will not necessarily be a complex structure, but we will see that it isn't hard to derive one from it. Thus as a first step, define a continuous map

$$\mathfrak{M}(E) \rightarrow \text{End}(E) : g \mapsto A$$

via the relation

$$\omega \equiv g(A\cdot, \cdot).$$

As is easy to check, the skew-symmetry of ω now implies that the fiberwise adjoint of A with respect to the bundle metric g is

$$A^* = -A,$$

so in particular A is a fiberwise *normal* operator, i.e. it commutes with its adjoint. Since A^*A is a positive definite symmetric form (again with respect to g), it has a

well defined square root, and there is thus a continuous map $\text{End}(E) \rightarrow \text{End}(E)$ that sends A to

$$J_g := A\sqrt{A^*A}^{-1}.$$

Now since A is normal, it also commutes with $\sqrt{A^*A}^{-1}$, and then $A^*A = -A^2$ implies $J_g^2 = -\mathbb{1}$. It is similarly straightforward to check that J_g is compatible with ω , and $J_g = J$ whenever $g = g_J$, hence the desired map is $\Phi(g) = J_g$. \square

At the beginning of Chapter 1, we showed that holomorphic curves in the standard \mathbb{C}^n have the important property that the area they trace out can be computed by integrating the standard symplectic structure. It is precisely this relation between symplectic structures and tame almost complex structures that makes the compactness theory of J -holomorphic curves possible. The computation generalizes as follows: assume $J \in \mathcal{J}^\tau(M, \omega)$ and let g_J be the Riemannian metric defined in (2.4). If $u : (\Sigma, j) \rightarrow (M, J)$ is a J -holomorphic curve and we choose holomorphic local coordinates (s, t) on a subset of Σ , then $\partial_t u = J\partial_s u$ implies that with respect to the metric g_J , $\partial_s u$ and $\partial_t u$ are orthogonal vectors of the same length. Thus the geometric area of the parallelogram spanned by these two vectors is simply

$$|\partial_s u|_{g_J} \cdot |\partial_t u|_{g_J} = |\partial_s u|_{g_J}^2 = \omega(\partial_s u, J\partial_s u) = \omega(\partial_s u, \partial_t u),$$

hence

$$(2.5) \quad \text{Area}_{g_J}(u) = \int_{\Sigma} u^* \omega.$$

Definition 2.11. For any $J \in \mathcal{J}^\tau(M, \omega)$, we define the **energy** of a J -holomorphic curve $u : (\Sigma, j) \rightarrow (M, J)$ by

$$E(u) = \int_{\Sigma} u^* \omega.$$

The following is an immediate consequence of (2.5).

Proposition 2.12. *If $J \in \mathcal{J}^\tau(M, \omega)$ then for every J -holomorphic curve $u : (\Sigma, j) \rightarrow (M, J)$, $E(u) \geq 0$, with equality if and only if u is locally constant.¹*

The energy as defined above is especially important in the case where the domain Σ is a closed surface. Then $u : \Sigma \rightarrow M$ represents a homology class $[u] := u_*[\Sigma] \in H_2(M)$, and the quantity $E(u)$ is not only nonnegative but also *topological*: it can be computed via the pairing $\langle [\omega], [u] \rangle$, and thus depends only on $[u] \in H_2(M)$ and $[\omega] \in H_{\text{dR}}^2(M)$. This implies an a priori energy bound for J -holomorphic curves in a fixed homology class, which we'll make considerable use of in applications.

For the next result, we can drop the assumption that M is a symplectic manifold, though the proof does make use of a (locally defined) symplectic structure. The result can be summarized by saying that for any reasonable moduli space of J -holomorphic curves, the *constant* curves form an open subset.

¹The term **locally constant** means that the restriction of u to each connected component of its domain Σ is constant.

Proposition 2.13. *Suppose Σ is a closed surface, $J_k \in \mathcal{J}(M)$ is a sequence of almost complex structures that converge in C^∞ to $J \in \mathcal{J}(M)$, and $u_k : (\Sigma, j_k) \rightarrow (M, J_k)$ is a sequence of non-constant pseudoholomorphic curves converging in C^∞ to a pseudoholomorphic curve $u : (\Sigma, j) \rightarrow (M, J)$. Then u is also not constant.*

PROOF. Assume u is constant and its image is $p \in M$. Choosing coordinates near p , we can assume without loss of generality that p is the origin in \mathbb{C}^n and u_k maps into a neighborhood of the origin, with almost complex structures J_k on \mathbb{C}^n converging to J such that $J(0) = i$. Then for sufficiently large k , the standard symplectic form ω_0 tames each J_k in a sufficiently small neighborhood of the origin, and $[u_k] = [u] = 0 \in H_2(M)$, implying $E(u_k) = \langle [\omega_0], [u_k] \rangle = 0$, thus u_k is also constant. \square

2.3. Linear Cauchy-Riemann type operators

Many important results about solutions to the nonlinear Cauchy-Riemann equation can be reduced to statements about solutions of corresponding linearized equations, thus it is important to understand the linearized equations first. Consider a Riemann surface (Σ, j) and a complex vector bundle $(E, J) \rightarrow (\Sigma, j)$ of (complex) rank n : this means that $E \rightarrow \Sigma$ is a real vector bundle of rank $2n$ and J is a complex structure on the bundle. We say that the bundle admits a **holomorphic structure** if Σ has an open covering $\{\mathcal{U}_\alpha\}$ with complex linear local trivializations $E|_{\mathcal{U}_\alpha} \rightarrow \mathcal{U}_\alpha \times \mathbb{C}^n$ whose transition maps are holomorphic functions from open subsets of Σ to $\mathrm{GL}(n, \mathbb{C})$.

On the space $C^\infty(\Sigma, \mathbb{C})$ of smooth complex valued functions, there are natural first-order differential operators

$$(2.6) \quad \bar{\partial} : f \mapsto df + i df \circ j$$

and

$$(2.7) \quad \partial : f \mapsto df - i df \circ j.$$

We can regard $\bar{\partial}$ as a linear map $C^\infty(\Sigma) \rightarrow \Gamma(\overline{\mathrm{Hom}}_{\mathbb{C}}(T\Sigma, \mathbb{C}))$, where the latter denotes the space of smooth sections of the bundle $\overline{\mathrm{Hom}}_{\mathbb{C}}(T\Sigma, \mathbb{C})$ of complex antilinear maps $T\Sigma \rightarrow \mathbb{C}$; similarly, ∂ maps $C^\infty(\Sigma)$ to $\Gamma(\mathrm{Hom}_{\mathbb{C}}(T\Sigma, \mathbb{C}))$.² Observe that the holomorphic functions $f : \Sigma \rightarrow \mathbb{C}$ are precisely those which satisfy $\bar{\partial}f \equiv 0$; the solutions of $\partial f \equiv 0$ are called **antiholomorphic**.

If $(E, J) \rightarrow (\Sigma, j)$ has a holomorphic structure, one can likewise define a natural operator on the space of sections $\Gamma(E)$,

$$\bar{\partial} : \Gamma(E) \rightarrow \Gamma(\overline{\mathrm{Hom}}_{\mathbb{C}}(T\Sigma, E)),$$

which is defined the same as (2.6) on any section written in a local holomorphic trivialization. We then call a section $v \in \Gamma(E)$ holomorphic if $\bar{\partial}v \equiv 0$, which is equivalent to the condition that it look holomorphic in all holomorphic local trivializations.

²Many authors prefer to write the spaces of sections of $\mathrm{Hom}_{\mathbb{C}}(T\Sigma, \mathbb{C})$ and $\overline{\mathrm{Hom}}_{\mathbb{C}}(T\Sigma, \mathbb{C})$ as $\Omega^{1,0}(\Sigma)$ and $\Omega^{0,1}(\Sigma)$ respectively, calling these sections “(1,0)-forms” and “(0,1)-forms.”

Exercise 2.14. Check that the above definition of $\bar{\partial} : \Gamma(E) \rightarrow \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$ doesn't depend on the trivialization if all transition maps are holomorphic. You may find Exercise 2.15 helpful. (Note that the operator $\partial f := df - i df \circ j$ is *not* similarly well defined on a bundle—it does depend on the trivialization in general.)

Exercise 2.15. Show that the $\bar{\partial}$ -operator on a holomorphic vector bundle satisfies the following Leibnitz identity: for any $v \in \Gamma(E)$ and $f \in C^\infty(\Sigma, \mathbb{C})$, $\bar{\partial}(fv) = (\bar{\partial}f)v + f(\bar{\partial}v)$.

Definition 2.16. A complex linear Cauchy-Riemann type operator on a complex vector bundle $(E, J) \rightarrow (\Sigma, j)$ is a complex linear map

$$D : \Gamma(E) \rightarrow \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$$

that satisfies the Leibnitz rule

$$(2.8) \quad D(fv) = (\bar{\partial}f)v + f(Dv)$$

for all $f \in C^\infty(\Sigma, \mathbb{C})$ and $v \in \Gamma(E)$.

One can think of this definition as analogous to the simplest modern definition of a connection on a vector bundle; in fact it turns out that every complex Cauchy-Riemann type operator is the complex linear part of some connection (see Exercises 2.18 and 2.19 below). The following is then the Cauchy-Riemann version of the existence of the Christoffel symbols.

Exercise 2.17.

- (1) Show that if D and D' are two complex linear Cauchy-Riemann type operators on a bundle $(E, J) \rightarrow (\Sigma, j)$, then there is a smooth complex linear bundle map $A : E \rightarrow \overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E)$ such that $D'v = Dv + Av$ for all $v \in \Gamma(E)$.
- (2) Show that in any local trivialization on a subset $\mathcal{U} \subset \Sigma$, every complex linear Cauchy-Riemann type operator D can be written in the form

$$Dv = \bar{\partial}v + Av,$$

for some smooth map $A : \mathcal{U} \rightarrow \text{End}_{\mathbb{C}}(\mathbb{C}^n)$.

Exercise 2.18. Show that if ∇ is any complex connection on E ,³ then $\nabla + J \circ \nabla \circ j$ is a complex linear Cauchy-Riemann operator.

Exercise 2.19. Prove the following converse of Exercise 2.18: if (E, J) is equipped with a Hermitian bundle metric and D is a complex linear Cauchy-Riemann type operator, then there is a unique Hermitian connection ∇ such that $D = \nabla + J \circ \nabla \circ j$.

Since connections exist in abundance on any vector bundle, there is always a Cauchy-Riemann type operator, even if (E, J) doesn't come equipped with a holomorphic structure. We now have the following analogue of Theorem 2.6 for bundles:

³By “complex connection” we mean that the parallel transport isomorphisms are complex linear. This is equivalent to the requirement that $\nabla : \Gamma(E) \rightarrow \Gamma(\text{Hom}_{\mathbb{R}}(T\Sigma, E))$ be a complex linear map.

Theorem 2.20. *For any complex linear Cauchy-Riemann type operator D on a complex vector bundle (E, J) over a Riemann surface (Σ, j) , there is a unique holomorphic structure on (E, J) such that the naturally induced $\bar{\partial}$ -operator is D .*

The proof can easily be reduced to the following local existence lemma, which is a special case of an analytical result that we'll prove in §2.6 (see Theorem 2.49):

Lemma 2.21. *Suppose D is a complex linear Cauchy-Riemann type operator on $(E, J) \rightarrow (\Sigma, j)$. Then for any $z \in \Sigma$ and $v_0 \in E_z$, there is a neighborhood $\mathcal{U} \subset \Sigma$ of z and a smooth section $v \in \Gamma(E|_{\mathcal{U}})$ such that $Dv = 0$ and $v(z) = v_0$.*

Exercise 2.22. Prove Theorem 2.20, assuming Lemma 2.21.

As we'll see in the next section, it's also quite useful to consider Cauchy-Riemann type operators that are only *real* linear, rather than complex.

Definition 2.23. A **real linear Cauchy-Riemann type operator** on a complex vector bundle $(E, J) \rightarrow (\Sigma, j)$ is a real linear map $D : \Gamma(E) \rightarrow \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$ such that (2.8) is satisfied for all $f \in C^\infty(\Sigma, \mathbb{R})$ and $v \in \Gamma(E)$.

The following is then merely an addendum to Exercise 2.17.

Exercise 2.24. Show that in any local trivialization on a subset $\mathcal{U} \subset \Sigma$, every real linear Cauchy-Riemann type operator D can be written in the form

$$Dv = \bar{\partial}v + Av,$$

for some smooth map $A : \mathcal{U} \rightarrow \text{End}_{\mathbb{R}}(\mathbb{C}^n)$, where $\text{End}_{\mathbb{R}}(\mathbb{C}^n)$ denotes the space of *real* linear maps on $\mathbb{C}^n = \mathbb{R}^{2n}$.

2.4. The linearization of $\bar{\partial}_J$ and critical points

We shall now see how linear Cauchy-Riemann type operators arise naturally from the nonlinear Cauchy-Riemann equation. Theorem 2.20 will then allow us already to prove something quite nontrivial: nonconstant J -holomorphic curves have only isolated critical points! It turns out that one can reduce this result to the corresponding statement about zeroes of holomorphic functions, a well known fact from complex analysis.

For the next few paragraphs, we will be doing a very informal version of “infinite dimensional differential geometry,” in which we assume that various spaces of smooth maps can sensibly be regarded as infinite dimensional smooth manifolds and vector bundles. For now this is purely for motivational purposes, thus we can avoid worrying about the technical details; when it comes time later to prove something using these ideas, we'll have to replace the spaces of smooth maps with Banach spaces, which will have to contain nonsmooth maps in order to attain completeness.

So, morally speaking, if (Σ, j) is a Riemann surface and (M, J) is an almost complex manifold, then the space of smooth maps $\mathcal{B} := C^\infty(\Sigma, M)$ is an infinite dimensional smooth manifold, and there is a vector bundle $\mathcal{E} \rightarrow \mathcal{B}$ whose fiber \mathcal{E}_u at $u \in \mathcal{B}$ is the space of smooth sections,

$$\mathcal{E}_u = \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, u^*TM)),$$

where we pull back J to define a complex bundle structure on $u^*TM \rightarrow \Sigma$. The tangent vectors at a point $u \in \mathcal{B}$ are simply vector fields along u , thus

$$T_u\mathcal{B} = \Gamma(u^*TM).$$

Now we define a section $\bar{\partial}_J : \mathcal{B} \rightarrow \mathcal{E}$ by

$$\bar{\partial}_J u = Tu + J \circ Tu \circ j.$$

This section is called the nonlinear Cauchy-Riemann operator, and its zeroes are precisely the J -holomorphic maps from Σ to M . Recall now that zero sets of smooth sections on bundles generically have a very nice structure—this follows from the implicit function theorem, of which we'll later use an infinite dimensional version. For motivational purposes only, we state here a finite dimensional version with geometric character. Recall that any section of a bundle can be regarded as an embedding of the base into the total space, thus we can always ask whether two sections are “transverse” when they intersect.

Theorem 2.25 (Finite dimensional implicit function theorem). *Suppose $E \rightarrow B$ is a smooth vector bundle of real rank k over an n -dimensional manifold and $s : B \rightarrow E$ is a smooth section that is everywhere transverse to the zero section. Then $s^{-1}(0) \subset B$ is a smooth submanifold of dimension $n - k$.*

The transversality assumption can easily be restated in terms of the *linearization* of the section s at a zero. The easiest way to define this is by choosing a connection ∇ on $E \rightarrow B$, as one can easily show that the linear map $\nabla s : T_p B \rightarrow E_p$ is independent of this choice at any point p where $s(p) = 0$; this follows from the fact that TE along the zero section has a canonical splitting into horizontal and vertical subspaces. Let us therefore denote the linearization at $p \in s^{-1}(0)$ by

$$Ds(p) : T_p B \rightarrow E_p.$$

Then the intersections of s with the zero section are precisely the set $s^{-1}(0)$, and these intersections are transverse if and only if $Ds(p)$ is a surjective map for all $p \in s^{-1}(0)$.

In later chapters we will devote considerable effort to finding ways of showing that the linearization of $\bar{\partial}_J$ at any $u \in \bar{\partial}_J^{-1}(0)$ is a surjective operator in the appropriate Banach space setting. With this as motivation, let us now deduce a formula for the linearization itself. It will be slightly easier to do this if we regard $\bar{\partial}_J$ as a section of the larger vector bundle $\widehat{\mathcal{E}}$ with fibers

$$\widehat{\mathcal{E}}_u = \Gamma(\text{Hom}_{\mathbb{R}}(T\Sigma, u^*TM)).$$

To choose a “connection” on $\widehat{\mathcal{E}}$, choose first a connection ∇ on M and assume that for any smoothly parametrized path $\tau \mapsto u_\tau \in \mathcal{B}$ and a section $\ell_\tau \in \widehat{\mathcal{E}}_{u_\tau} = \Gamma(\text{Hom}_{\mathbb{R}}(T\Sigma, u_\tau^*TM))$ along the path, the covariant derivative $\nabla_\tau \ell_\tau \in \widehat{\mathcal{E}}_{u_\tau}$ should take the form

$$(\nabla_\tau \ell_\tau)X = \nabla_\tau (\ell_\tau(X)) \in (u^*TM)_z = T_{u(z)}M$$

for $z \in \Sigma$, $X \in T_z \Sigma$. Then $\nabla_\tau \ell_\tau$ doesn't depend on the choice of ∇ at any value of τ for which $\ell_\tau = 0$.

Now given $u \in \bar{\partial}_J^{-1}(0)$, consider a smooth family of maps $\{u_\tau\}_{\tau \in (-1,1)}$ with $u_0 = u$, and write $\partial_\tau u_\tau|_{\tau=0} =: \eta \in \Gamma(u^*TM)$. By definition, the linearization

$$D\bar{\partial}_J(u) : \Gamma(u^*TM) \rightarrow \Gamma(\text{Hom}_{\mathbb{R}}(T\Sigma, u^*TM))$$

will be the unique linear map such that

$$D\bar{\partial}_J(u)\eta = \nabla_\tau (\bar{\partial}_J u_\tau)|_{\tau=0} = \nabla_\tau [Tu_\tau + J(u_\tau) \circ Tu_\tau \circ j]|_{\tau=0}.$$

To simplify this expression, choose holomorphic local coordinates $s + it$ near the point $z \in \Sigma$ and consider the action of the above expression on the vector ∂_s : this gives

$$\nabla_\tau [\partial_s u_\tau + J(u_\tau) \partial_t u_\tau]|_{\tau=0}.$$

The expression simplifies further if we assume ∇ is a *symmetric* connection on M ; this is allowed since the end result will not depend on the choice of connection. In this case $\nabla_\tau \partial_s u_\tau|_{\tau=0} = \nabla_s \partial_\tau u_\tau|_{\tau=0} = \nabla_s \eta$ and similarly for the derivative by t , thus the above becomes

$$\nabla_s \eta + J(u) \nabla_t \eta + (\nabla_\eta J) \partial_t u.$$

Taking the coordinates back out, we're led to the following expression for the linearization of $\bar{\partial}_J$:

$$(2.9) \quad D\bar{\partial}_J(u)\eta = \nabla \eta + J(u) \circ \nabla \eta \circ j + (\nabla_\eta J) \circ Tu \circ j.$$

Though we shouldn't have expected it a priori since we enlarged \mathcal{E} to $\widehat{\mathcal{E}}$, it turns out that the right hand side of (2.9) does belong to \mathcal{E}_u , i.e. it is a complex antilinear bundle map $T\Sigma \rightarrow u^*TM$.

Exercise 2.26. Verify that if $u \in \bar{\partial}_J^{-1}(0)$, then for any $\eta \in \Gamma(u^*TM)$, the bundle map $T\Sigma \rightarrow u^*TM$ defined by the right hand side of (2.9) is complex antilinear. *Hint: show first that $\nabla_X J$ always anticommutes with J for any vector X .*

To move back into the realm of solid mathematics, let us now regard (2.9) as a definition, i.e. to any smooth J -holomorphic map $u : \Sigma \rightarrow M$ we associate the operator

$$\mathbf{D}_u := D\bar{\partial}_J(u),$$

which is a real linear map taking sections of u^*TM to sections of $\overline{\text{Hom}_{\mathbb{C}}}(T\Sigma, u^*TM)$. The following exercise is straightforward but important.

Exercise 2.27. Show that \mathbf{D}_u is a real linear Cauchy-Riemann type operator on u^*TM .

With this and Theorem 2.20 to work with, it is already quite easy to prove that J -holomorphic curves have isolated critical points. The key idea, due to Ivashkovich and Shevchishin [IS99], is to use the linearized operator \mathbf{D}_u to define a holomorphic structure on $\text{Hom}_{\mathbb{C}}(T\Sigma, u^*TM)$ so that du becomes a holomorphic section. Observe first that since (Σ, j) is a complex manifold, the bundle $T\Sigma \rightarrow \Sigma$ has a natural holomorphic structure, so one can speak of holomorphic vector fields on Σ . In general such vector fields will be defined only locally, but this is sufficient for our purposes.

Exercise 2.28. A map $\varphi : (\Sigma, j) \rightarrow (\Sigma, j)$ is holomorphic if and only if it satisfies the low dimensional case of the nonlinear Cauchy-Riemann equation, $\bar{\partial}_j \varphi = 0$. The simplest example of such a map is the identity $\text{Id} : \Sigma \rightarrow \Sigma$, and the linearization \mathbf{D}_{Id} gives an operator $\Gamma(T\Sigma) \rightarrow \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, T\Sigma))$. Show that \mathbf{D}_{Id} is complex linear, and in fact it is the natural Cauchy-Riemann operator determined by the holomorphic structure of $T\Sigma$. *Hint: in holomorphic local coordinates this is almost obvious.*

Lemma 2.29. Suppose X is a holomorphic vector field on some open subset $\mathcal{U} \subset \Sigma$, $\mathcal{U}' \subset \mathcal{U}$ is another open subset and $\epsilon > 0$ a number such that the flow $\varphi_X^t : \mathcal{U}' \rightarrow \Sigma$ is well defined for all $t \in (-\epsilon, \epsilon)$. Then the maps φ_X^t are holomorphic.

PROOF. Working in local holomorphic coordinates, this reduces to the following claim: if $\mathcal{U} \subset \mathbb{C}$ is an open subset containing a smaller open set $\mathcal{U}' \subset \mathcal{U}$, $X : \mathcal{U} \rightarrow \mathbb{C}$ is a holomorphic function and $\varphi^\tau : \mathcal{U}' \rightarrow \mathbb{C}$ satisfies

$$(2.10) \quad \begin{aligned} \partial_\tau \varphi^\tau(z) &= X(\varphi^\tau(z)), \\ \varphi^0(z) &= z \end{aligned}$$

for $\tau \in (-\epsilon, \epsilon)$, then φ^τ is holomorphic for every τ . To see this, apply the operator $\bar{\partial} := \partial_s + i\partial_t$ to both sides of (2.10) and exchange the order of partial derivatives: this gives

$$\frac{\partial}{\partial \tau} \bar{\partial} \varphi^\tau(z) = X'(\varphi^\tau(z)) \cdot \bar{\partial} \varphi^\tau(z).$$

For any fixed $z \in \mathcal{U}'$, this is a linear differential equation for the complex valued path $\tau \mapsto \bar{\partial} \varphi^\tau(z)$. Since it begins at zero, uniqueness of solutions implies that it is identically zero. \square

Lemma 2.30. For any holomorphic vector field X defined on an open subset $\mathcal{U} \subset \Sigma$, $\mathbf{D}_u[Tu(X)] = 0$ on \mathcal{U} .

PROOF. By shrinking \mathcal{U} if necessary, we can assume that the flow $\varphi_X^t : \mathcal{U} \rightarrow \Sigma$ is well defined for sufficiently small $|t|$, and by Lemma 2.29 it is holomorphic, hence the maps $u \circ \varphi_X^t$ are also J -holomorphic. Then $\bar{\partial}_J(u \circ \varphi_X^t) = 0$ and

$$\mathbf{D}_u[Tu(X)] = \nabla_t [\bar{\partial}_J(u \circ \varphi_X^t)]|_{t=0} = 0.$$

\square

The Cauchy-Riemann type operator \mathbf{D}_u is real linear, but one can easily define a complex linear operator by projecting out the antilinear part:

$$\mathbf{D}_u^{\mathbb{C}} = \frac{1}{2} (\mathbf{D}_u - J \circ \mathbf{D}_u \circ J).$$

This defines a complex linear map $\Gamma(u^*TM) \rightarrow \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, u^*TM))$.

Exercise 2.31. Show that $\mathbf{D}_u^{\mathbb{C}}$ is a complex linear Cauchy-Riemann type operator.

In light of Exercise 2.31 and Theorem 2.20, the induced bundle $u^*TM \rightarrow \Sigma$ for any smooth J -holomorphic curve $u : \Sigma \rightarrow M$ admits a holomorphic structure for

which holomorphic sections satisfy $\mathbf{D}_u^{\mathbb{C}}\eta = 0$. Moreover, Lemma 2.30 implies that for any local holomorphic vector field X on Σ ,

$$\mathbf{D}_u^{\mathbb{C}}[Tu(X)] = \frac{1}{2}\mathbf{D}_u[Tu(X)] - \frac{1}{2}J\mathbf{D}_u[J \circ Tu(X)] = \frac{1}{2}J\mathbf{D}_u[Tu(jX)] = 0,$$

where we've used the nonlinear Cauchy-Riemann equation for u and the fact that jX is also holomorphic. Thus $Tu(X)$ is a holomorphic section on u^*TM whenever X is holomorphic on $T\Sigma$. Put another way, the holomorphic bundle structures on $T\Sigma$ and u^*TM naturally induce a holomorphic structure on $\mathrm{Hom}_{\mathbb{C}}(T\Sigma, u^*TM)$, and the section $du \in \Gamma(\mathrm{Hom}_{\mathbb{C}}(T\Sigma, u^*TM))$ is then holomorphic. We've proved:

Theorem 2.32. *For any smooth J -holomorphic map $u : \Sigma \rightarrow M$, the complex linear part of the linearization \mathbf{D}_u induces on $\mathrm{Hom}_{\mathbb{C}}(T\Sigma, u^*TM)$ a holomorphic structure such that du is a holomorphic section.*

Corollary 2.33. *If $u : \Sigma \rightarrow M$ is smooth, J -holomorphic and not constant, then the set $\mathrm{Crit}(u) := \{z \in \Sigma \mid du(z) = 0\}$ is discrete.*

Actually we've proved more: using a holomorphic trivialization of the bundle $\mathrm{Hom}_{\mathbb{C}}(T\Sigma, u^*TM)$ near any $z_0 \in \mathrm{Crit}(u)$, one can choose holomorphic coordinates identifying z_0 with $0 \in \mathbb{C}$ and write $du(z)$ in the trivialization as

$$du(z) = z^k F(z),$$

where $k \in \mathbb{N}$ and F is a nonzero \mathbb{C}^n -valued holomorphic function. This means that each critical point of u has a well defined and positive *order* (the number k), as well as a *tangent plane* (the complex 1-dimensional subspace spanned by $F(0)$ in the trivialization). We will see this again when we investigate intersections in §2.11, and it will also prove useful later when we discuss “automatic” transversality.

Remark 2.34. The above results for the critical set of a J -holomorphic curve u remain valid if we don't require smoothness but only assume $J \in C^1$ and $u \in C^2$: then u^*TM and $\overline{\mathrm{Hom}_{\mathbb{C}}(T\Sigma, u^*TM)}$ are complex vector bundles of class C^1 and du is a C^1 -section, but turns out to be holomorphic with respect to a system of non-smooth trivializations which have holomorphic (and therefore smooth!) transition functions. One can prove this using the weak regularity assumptions in Theorem 2.49 below; in practice of course, the regularity results of §2.10 will usually allow us to avoid such questions altogether.

2.5. Linear elliptic regularity

Until now we've usually assumed that our J -holomorphic maps $u : \Sigma \rightarrow M$ are smooth, but for technical reasons we'll later want to allow maps with weaker, Sobolev-type regularity assumptions. In the end it all comes to the same thing, because if J is smooth, then it turns out that all J -holomorphic curves are also smooth. In the integrable case, one can choose coordinates in M so that $J = i$ and J -holomorphic curves are honestly holomorphic, then this smoothness statement is a well known corollary of the Cauchy integral formula. The nonintegrable case requires more work and makes heavy use of the machinery of elliptic PDE theory. We will not cover this subject in full detail—in particular, a few estimates will

have to be taken as black boxes—but we shall give an overview of the regularity results that we'll need and try to explain why they're true.⁴ As consequences, in this section we will see how to prove smoothness of solutions to linear Cauchy-Riemann type equations and also derive an important surjectivity property of the $\bar{\partial}$ -operator, which will later help in proving local existence results. The discussion necessarily begins with the *linear* case, and we will address the nonlinear case in §2.8. It should also be mentioned that the estimates in this section have more than just local consequences: they will be crucial later when we discuss the global Fredholm and compactness theory of J -holomorphic curves.

Let us first look at a much simpler differential equation to illustrate the idea of elliptic regularity. Suppose $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a function of class C^k and we have a C^1 solution to the nonlinear ODE,

$$(2.11) \quad \dot{x} = F(x).$$

Then if $k \geq 1$, the right hand side is clearly of class C^1 , thus so is \dot{x} , implying that x is actually C^2 . If $k \geq 2$, we can repeat the argument and find that x is C^3 and so on; in the end we find $x \in C^{k+1}$, i.e. x is at least one step smoother than F . This induction is the simplest example of an “elliptic bootstrapping argument”.

The above argument is extremely easy because the left hand side of (2.11) tells us everything we'd ever want to know about the first derivative of our solution. The situation for a first-order PDE is no longer so simple: e.g. consider the usual Cauchy-Riemann operator for functions $\mathbb{C} \rightarrow \mathbb{C}$,

$$\bar{\partial} = \partial_s + i\partial_t,$$

and the associated linear inhomogeneous equation

$$\bar{\partial}u = f.$$

Now the left hand side carries *part*, but not *all* of the information one could want to have about du : one can say that $\partial_s u + i\partial_t u$ is at least as smooth as f , but this doesn't immediately imply the same statement for each of $\partial_s u$ and $\partial_t u$. What we need is a way to estimate du (in some suitable norm) in terms of u and $\bar{\partial}u$, and this turns out to be possible precisely because $\bar{\partial}$ is an *elliptic* operator. We will not attempt here to define precisely what “elliptic” means;⁵ in practice, a differential operator is called elliptic if it can be shown to satisfy a fundamental estimate of the type stated in Theorem 2.35 below.

We briefly recall some ideas and notation from the theory of Sobolev spaces (see e.g. [Eva98] or the appendices of [MS04]). If $k \in \mathbb{N}$, $p \geq 1$ and $\mathcal{U} \subset \mathbb{C}$ is an open subset, then $W^{0,p}(\mathcal{U}) := L^p(\mathcal{U})$ is the space of (real or complex valued) functions of class L^p on \mathcal{U} , and inductively $W^{k,p}(\mathcal{U})$ denotes the space of functions in $L^p(\mathcal{U})$ that have weak derivatives in $W^{k-1,p}(\mathcal{U})$. Recall that a locally integrable function

⁴For a more comprehensive treatment from a slightly different perspective, Appendix B of [MS04] is indispensable.

⁵The definition of a second-order elliptic operator is treated in most introductory texts on PDE theory, see for example [Eva98]. A general definition for all orders may be found in [DN55].

$u \in L^1_{\text{loc}}(\mathcal{U})$ is said to have **weak derivative** $\partial_s u = g \in L^1_{\text{loc}}(\mathcal{U})$ if for every smooth compactly supported function $\varphi \in C_0^\infty(\mathcal{U})$,

$$\int_{\mathcal{U}} \varphi g = - \int_{\mathcal{U}} (\partial_s \varphi) u.$$

In other words, $g = \partial_s u$ in the sense of distributions. One defines $\partial_t u$ and higher order weak derivatives via similar formulas based on integration by parts, and in this way one can also speak of **weak solutions** to equations such as $\bar{\partial} u = f$, where in general u and f need not be more regular than L^1_{loc} (or even more generally, distributions). We say that a function on \mathcal{U} is of class $W^{k,p}_{\text{loc}}$ if it is in $W^{k,p}(\mathcal{U}')$ for every open subset \mathcal{U}' with compact closure $\bar{\mathcal{U}}' \subset \mathcal{U}$. In the following, we will consider Sobolev spaces of maps valued in complex vector spaces such as \mathbb{C}^n ; we'll specify the target space by writing e.g. $W^{k,p}(\mathcal{U}, \mathbb{C}^n)$ whenever there is danger of confusion. The symbols B and B_r will denote the open balls in \mathbb{C} of radius 1 and r respectively. We will often make use of the *Sobolev embedding theorem*, which in the present context implies that if $\mathcal{U} \subset \mathbb{C}$ is a bounded open domain with smooth boundary and $kp > 2$, then there are natural continuous inclusions

$$W^{k+d,p}(\mathcal{U}) \hookrightarrow C^d(\mathcal{U})$$

for each integer $d \geq 0$. In fact these inclusions are compact linear operators, as are the obvious inclusions

$$W^{k,p}(\mathcal{U}) \hookrightarrow W^{k-1,p}(\mathcal{U}).$$

Theorem 2.35 (Calderón-Zygmund inequality). *For each $p \in (1, \infty)$, there is a constant $c > 0$ such that for every $u \in C_0^\infty(B, \mathbb{C}^n)$,*

$$\|u\|_{W^{1,p}} \leq c \|\bar{\partial} u\|_{L^p}.$$

Exercise 2.36. Assuming the theorem above, differentiate the equation $\bar{\partial} u = f$ and argue by induction to prove the following generalization: for each $k \in \mathbb{N}$ and $p \in (1, \infty)$ there is a constant $c > 0$ such that

$$\|u\|_{W^{k,p}} \leq c \|\bar{\partial} u\|_{W^{k-1,p}}$$

for all $u \in C_0^\infty(B)$. By a density argument, show that this also holds for all $u \in W_0^{k,p}(B)$, where the latter denotes the closure of $C_0^\infty(B)$ in $W^{k,p}(B)$.

Theorem 2.35 follows from estimates of certain singular integrals carried out in [CZ52, CZ56] (see also the discussion in Appendix B of [MS04], which expresses it in terms of the Laplace operator). More general versions for elliptic systems of all orders appear in [DN55], and versions with boundary conditions are treated in [ADN59, ADN64]. Before launching into a cursory discussion of the proof, let us see how this estimate can be used to prove a basic local regularity result for the linear inhomogeneous Cauchy-Riemann equation. We will later improve this result to apply to weak solutions of class L^1_{loc} (see Theorem 2.45 below).

Proposition 2.37. *Suppose $u \in W^{1,p}(B, \mathbb{C}^n)$ and $\bar{\partial} u \in W^{k,p}(B, \mathbb{C}^n)$ for some $p \in (1, \infty)$. Then $u \in W^{k+1,p}(B_r, \mathbb{C}^n)$ for any $r < 1$, and there is a constant c , depending on r and p but not on u , such that*

$$(2.12) \quad \|u\|_{W^{k+1,p}(B_r)} \leq c \|u\|_{W^{k,p}(B)} + c \|\bar{\partial} u\|_{W^{k,p}(B)}.$$

Corollary 2.38. *If $f : B \rightarrow \mathbb{C}^n$ is smooth, then every solution to $\bar{\partial}u = f$ of class $W^{1,p}$ for some $p \in (1, \infty)$ is also smooth. Moreover, given sequences $f_k \rightarrow f$ converging in $C^\infty(B)$ and $u_k \rightarrow u$ converging in $W^{1,p}(B)$ and satisfying $\bar{\partial}u_k = f_k$, the sequence u_k also converges in C_{loc}^∞ on B .*

Exercise 2.39. Prove the corollary.

PROOF OF PROP. 2.37. Write $\bar{\partial}u = f$. It will suffice to consider the case $k = 1$, as once this is settled, the result follows from an easy induction argument using the fact that any derivative $D^\alpha u$ of u satisfies $\bar{\partial}D^\alpha u = D^\alpha f$. (Here α is a multiindex, so D^α may be any differential operator of order one or greater.)

Now assuming $u, f \in W^{1,p}(B)$, we'd first like to prove that u is of class $W^{2,p}$ on B_r for any $r < 1$. The idea is to show that $\partial_s u$ (and similarly $\partial_t u$) is of class $W^{1,p}$ by expressing it as a limit of the difference quotients,

$$u^h(s, t) := \frac{u(s + h, t) - u(s, t)}{h}$$

as $h > 0$ shrinks to zero. These functions are clearly well defined and belong to $W^{1,p}(B_r)$ if h is sufficiently small, and it is straightforward (e.g. using approximation by smooth functions) to show that $u^h \rightarrow \partial_s u$ in $L^p(B_r)$ as $h \rightarrow 0$. The significance of Theorem 2.35 is that it gives us a uniform $W^{1,p}$ -bound on u^h with respect to h . Indeed, pick a cutoff function $\beta \in C_0^\infty(B)$ that equals 1 on B_r . Then $\beta u^h \in W_0^{1,p}(B)$ and thus satisfies the estimate of Theorem 2.35 (cf. Exercise 2.36). We compute

$$\begin{aligned} (2.13) \quad \|u^h\|_{W^{1,p}(B_r)} &\leq \|\beta u^h\|_{W^{1,p}(B)} \leq c\|\bar{\partial}(\beta u^h)\|_{L^p(B)} \\ &= c\|(\bar{\partial}\beta)u^h + \beta(\bar{\partial}u^h)\|_{L^p(B)} \leq c'\|u^h\|_{L^p(B)} + c'\|f^h\|_{L^p(B)}, \end{aligned}$$

and observe that the right hand side is bounded as $h \rightarrow 0$ because $u^h \rightarrow \partial_s u$ and $f^h \rightarrow \partial_s f$ in L^p .

In light of this bound, the Banach-Alaoglu theorem implies that any sequence u^{h_k} with $h_k \rightarrow 0$ has a weakly convergent subsequence in $W^{1,p}(B_r)$. But since u^h already converges to $\partial_s u$ in $L^p(B_r)$, the latter must also be the weak $W^{1,p}$ -limit, implying $\partial_s u \in W^{1,p}(B_r)$. Now the estimate (2.12) follows from (2.13) by observing,

$$\begin{aligned} \|\partial_s u\|_{W^{1,p}(B_r)} &\leq \limsup_{h \rightarrow 0} \|u^h\|_{W^{1,p}(B_r)} \\ &\leq c\|\partial_s u\|_{L^p(B)} + c\|\partial_s f\|_{L^p(B)} \leq c\|u\|_{W^{1,p}(B)} + c\|f\|_{W^{1,p}(B)}. \end{aligned}$$

□

Exercise 2.40. Use the above result to show that for any real linear Cauchy-Riemann type operator D on a vector bundle $(E, J) \rightarrow (\Sigma, j)$, continuously differentiable solutions of $D\eta = 0$ are always smooth. *Note: due to Exercise 2.24, this reduces to showing that solutions $u \in W^{1,p}(B, \mathbb{C}^n)$ of $(\bar{\partial} + A)u = 0$ are smooth if $A : B \rightarrow \text{End}_{\mathbb{R}}(\mathbb{C}^n)$ is smooth.*

The proof of Theorem 2.35 contains several easy steps and one that is hard. For the easy part: observe first that by the Poincaré inequality, $\|u\|_{L^p}$ can be bounded in

terms of $\|du\|_{L^p}$ for any $u \in C_0^\infty(B)$, thus it will suffice to bound the first derivatives in terms of $\bar{\partial}u$. For this it is natural to consider the conjugate of the $\bar{\partial}$ -operator,

$$\partial := \partial_s - i\partial_t,$$

as the expressions $\bar{\partial}u$ and ∂u together can produce both $\partial_s u$ and $\partial_t u$ by linear combinations. Thus we are done if we can show that $\|\partial u\|_{L^p}$ is bounded in terms of $\|\bar{\partial}u\|_{L^p}$.

It is easy to see why this is true in the case $p = 2$: the following “symplectic” argument is borrowed from [Sik94]. It suffices to set $n = 1$ and consider compactly supported smooth functions $u : B \rightarrow \mathbb{C}$. Using the coordinate $z = s + it$, define the differential operators

$$\partial_z = \frac{\partial}{\partial z} = \frac{1}{2}(\partial_s - i\partial_t) \quad \partial_{\bar{z}} = \frac{\partial}{\partial \bar{z}} = \frac{1}{2}(\partial_s + i\partial_t)$$

and corresponding complex valued 1-forms

$$dz = d(s + it) = ds + i dt \quad d\bar{z} = d(s - it) = ds - i dt.$$

Observe that ∂_z and $\partial_{\bar{z}}$ are the same as ∂ and $\bar{\partial}$ respectively up to a factor of two,⁶ and we now have $du = \partial_z u dz + \partial_{\bar{z}} u d\bar{z}$. Viewing \mathbb{C} as a subset of the Riemann sphere S^2 , $u \in C^\infty(B, \mathbb{C})$ naturally extends to a smooth function $S^2 \rightarrow \mathbb{C}$, and $d(u d\bar{u}) = du \wedge d\bar{u}$ defines an exact complex valued 2-form with support in B . Now by Stokes’ theorem,

$$\begin{aligned} 0 &= \int_{S^2} du \wedge d\bar{u} = \int_{S^2} (\partial_z u dz + \partial_{\bar{z}} u d\bar{z}) \wedge (\partial_z \bar{u} dz + \partial_{\bar{z}} \bar{u} d\bar{z}) \\ &= \frac{1}{4} \int_{S^2} (|\partial u|^2 - |\bar{\partial} u|^2) dz \wedge d\bar{z}, \end{aligned}$$

hence $\|\partial u\|_{L^2} = \|\bar{\partial} u\|_{L^2}$.

For $p \neq 2$, a bound on $\|\partial u\|_{L^p}$ can be found by rephrasing the equation $\bar{\partial}u = f$ in terms of fundamental solutions. We recall the basic idea: a fundamental solution to the $\bar{\partial}$ -equation is an L^1_{loc} -function $K : \mathbb{C} \rightarrow \mathbb{C}$ that satisfies $\bar{\partial}K = \delta$ in the sense of distributions, where δ is the Dirac delta function, i.e. the distribution whose action on any test function $\varphi \in C_0^\infty$ is $\delta(\varphi) = \varphi(0)$. Then weak solutions of $\bar{\partial}u = f$ can be expressed as convolutions $u = K * f$, since $\bar{\partial}(K * f) = \bar{\partial}K * f = \delta * f = f$. To make all this precise, define the function $K \in L^1_{\text{loc}}(\mathbb{C}, \mathbb{C})$ to be

$$K(z) = \frac{1}{2\pi z},$$

and for any $f \in C_0^\infty(\mathbb{C}, \mathbb{C}^n)$, define $Tf : \mathbb{C} \rightarrow \mathbb{C}^n$ to be the convolution

$$(2.14) \quad Tf(z) = K * f(z) = \int_{\mathbb{C}} K(z - \zeta) f(\zeta) \frac{d\zeta \wedge d\bar{\zeta}}{-2i} = \frac{1}{2\pi} \int_{\mathbb{C}} \frac{f(\zeta)}{z - \zeta} \frac{d\zeta \wedge d\bar{\zeta}}{-2i},$$

where we use the notation $\frac{d\zeta \wedge d\bar{\zeta}}{-2i}$ to abbreviate the standard Lebesgue measure on \mathbb{C} with complex coordinate ζ . This integral is clearly well defined if $f \in C_0^\infty$, and it is not too hard to show that it gives a smooth function $Tf \in C^\infty(\mathbb{C}, \mathbb{C}^n)$ which decays

⁶It is in some sense more natural to define the operators $\bar{\partial}$ and ∂ with the factor of $1/2$ included, but we have dropped this for the sake of notational convenience.

to zero at infinity and satisfies $\bar{\partial}Tf = f$ (see for example [HZ94, Appendix A4]). It follows that for $u \in C_0^\infty(\mathbb{C}, \mathbb{C}^n)$, $\bar{\partial}u = f$ if and only if $u = Tf$. Indeed, if u has compact support then so does $f = \bar{\partial}u$, thus Tf is well defined and we have $\bar{\partial}(u - Tf) = 0$. But since Tf decays at infinity, this means $u - Tf$ is a bounded holomorphic function on \mathbb{C} that approaches zero at infinity, hence it is identically zero.

By the above remarks, it suffices to show that the function

$$\Pi f := \partial T f$$

satisfies an L^p -bound in terms of $\|f\|_{L^p}$ whenever $f \in C_0^\infty(B)$. Naively, one can derive an expression for $\partial T f = \partial(K * f)$ by differentiating the fundamental solution: it should be the convolution of f with

$$\partial K(z) = 2 \frac{\partial}{\partial z} \frac{1}{2\pi z} = -\frac{1}{\pi z^2}.$$

Thus we are led to the expression

$$\Pi f(z) = \lim_{\epsilon \rightarrow 0} \int_{|\zeta - z| \geq \epsilon} -\frac{f(\zeta)}{\pi(z - \zeta)^2} \frac{d\zeta \wedge d\bar{\zeta}}{-2i},$$

where we've attached a limit in order to define the integral since $\partial K \notin L_{\text{loc}}^1$; this is a *Cauchy principal value* integral. One can now verify explicitly that this limit is well defined if $f \in C_0^\infty$, and it computes $\partial T f(z)$. The hard work is thus reduced to the following estimates.

Lemma 2.41. *For each $p \in (1, \infty)$ there is a constant $c > 0$ such that for every $f \in C_0^\infty(B)$,*

$$\|Tf\|_{L^p} \leq c\|f\|_{L^p} \quad \text{and} \quad \|\Pi f\|_{L^p} \leq c\|f\|_{L^p}.$$

This is quite hard work indeed; complete proofs may be found in [Vek62, Part I, §5–9], or in a slightly different context, [MS04, Appendix B]. Taking this as a black box, the lemma implies Theorem 2.35.

In fact, Lemma 2.41 implies a stronger result about the Cauchy-Riemann operator that we will find useful. Since $C_0^\infty(B)$ is dense in $L^p(B)$, the first inequality in the lemma says that T can be extended to a bounded linear operator on $L^p(B)$, but the second inequality implies even more: we already have a bound on $\bar{\partial}Tf$ since this is simply f , so a bound on $\Pi f = \partial T f$ gives a bound on the whole first derivative of Tf , and we conclude:

Proposition 2.42. *For each $p \in (1, \infty)$, the operator T of (2.14) extends to a bounded linear operator $T : L^p(B) \rightarrow W^{1,p}(B)$, which is a right inverse of $\bar{\partial} : W^{1,p}(B) \rightarrow L^p(B)$.*

The upshot is that the equation $\bar{\partial}u = f$ can be solved for any $f \in L^p(B)$, and in a way that controls the first derivatives of the solution. This can be improved further using the previous regularity results:

Theorem 2.43. *For any integer $k \geq 0$ and $p \in (1, \infty)$, the operator $\bar{\partial} : W^{k+1,p}(B) \rightarrow W^{k,p}(B)$ admits a bounded right inverse*

$$\widehat{T} : W^{k,p}(B) \rightarrow W^{k+1,p}(B),$$

i.e. $\bar{\partial}\widehat{T}f = f$ for all $f \in W^{k,p}(B)$.

PROOF. Prop. 2.42 proves the result for $k = 0$, so we proceed by induction, assuming the result is proven already for $k - 1$. Pick $R > 1$, and for each ℓ let

$$W^{\ell,p}(B) \rightarrow W^{\ell,p}(B_R) : f \mapsto \hat{f}$$

denote a bounded linear *extension* operator, i.e. \hat{f} satisfies $\hat{f}|_B = f$ and $\|\hat{f}\|_{W^{\ell,p}(B_R)} \leq c\|f\|_{W^{\ell,p}(B)}$ for some $c > 0$. Then by assumption there is a bounded operator

$$T_R : W^{k-1,p}(B_R) \rightarrow W^{k,p}(B_R)$$

that is a right inverse of $\bar{\partial}$, hence $u := T_R \hat{f}$ satisfies $\bar{\partial}u = \hat{f}$. But then if $\hat{f} \in W^{k,p}(B_R)$, Prop. 2.37 implies that $u \in W^{k+1,p}(B)$ and

$$\begin{aligned} \|u\|_{W^{k+1,p}(B)} &\leq c\|u\|_{W^{k,p}(B_R)} + c\|\hat{f}\|_{W^{k,p}(B_R)} \\ &\leq \|T_R\| \cdot \|\hat{f}\|_{W^{k-1,p}(B_R)} + c_1\|f\|_{W^{k,p}(B)} \leq c_2\|f\|_{W^{k,p}(B)}. \end{aligned}$$

□

Now that we are guaranteed to have nice solutions of the equation $\bar{\partial}u = f$, we can also improve the previous regularity results to apply to more general weak solutions. We begin with the simple fact that “weakly” holomorphic functions are actually smooth.

Lemma 2.44. *If $u \in L^1(B)$ is a weak solution of $\bar{\partial}u = 0$, then u is smooth.*

PROOF. By taking real and imaginary parts, it suffices to prove the same statement for real valued weak solutions of the Laplace equation: thus consider a function $u \in L^1(B, \mathbb{R})$ such that $\Delta u = 0$ in the sense of distributions. On B_r for any $r < 1$ we can approximate u by smooth functions u_ϵ using a standard mollifier,

$$u_\epsilon = j_\epsilon * u,$$

so that $u_\epsilon \rightarrow u$ in $L^1(B_r)$ as $\epsilon \rightarrow 0$. Moreover, $\Delta u_\epsilon = j_\epsilon * \Delta u = 0$, thus the u_ϵ are harmonic. This implies that they satisfy the mean value property, so for every sufficiently small ball $B_\delta(z)$ about any point $z \in B_r$,

$$u_\epsilon(z) = \frac{1}{\pi\delta^2} \int_{B_\delta(z)} u_\epsilon(s, t) ds dt.$$

By L^1 -convergence, this expression converges pointwise in a neighborhood of z to the map

$$z \mapsto \frac{1}{\pi\delta^2} \int_{B_\delta(z)} u(s, t) ds dt.$$

The latter is continuous, and must be equal to u almost everywhere, thus u satisfies the mean value property and is therefore a smooth harmonic function (see [Eva98, §2.2.3]). □

Theorem 2.45. *Suppose $f \in W^{k,p}(B, \mathbb{C}^n)$ for some $p \in (1, \infty)$ and $u \in L^1(B, \mathbb{C}^n)$ is a weak solution of the equation $\bar{\partial}u = f$. Then $u \in W^{k+1,p}(B_r, \mathbb{C}^n)$ for any $r < 1$.*

PROOF. By Theorem 2.43, there is a solution $\eta \in W^{k+1,p}(B)$ to $\bar{\partial}\eta = f$, and then $\bar{\partial}(u - \eta) = 0$. Lemma 2.44 then implies that $u - \eta$ is smooth and hence in $W^{k+1,p}(B_r)$ for all $r < 1$, thus u is also in $W^{k+1,p}(B_r)$. \square

Corollary 2.46. *Suppose $1 < p < \infty$, k is a nonnegative integer, $A \in L^\infty(B, \text{End}_{\mathbb{R}}(\mathbb{C}^n))$, $f \in W^{k,p}(B, \mathbb{C}^n)$ and $u \in L^p(B, \mathbb{C}^n)$ is a weak solution of the equation $\bar{\partial}u + Au = f$. Then $u \in W^{1,p}(B_r, \mathbb{C}^n)$ for any $r < 1$. Moreover if A is smooth, then $u \in W^{k+1,p}(B_r, \mathbb{C}^n)$, and in particular u is smooth if f is smooth.*

PROOF. We have $\bar{\partial}u = -Au + f$ of class L^p , thus $u \in W^{1,p}(B_r)$ by Theorem 2.45. If A is also smooth and $k \geq 1$, then $-Au + f$ is now of class $W^{1,p}$, so $u \in W^{2,p}(B_r)$, and repeating this argument inductively, we eventually find $u \in W^{k+1,p}(B_r)$. \square

The invertibility results for $\bar{\partial}$ will also be useful for proving more general local existence results, because the property of having a bounded right inverse is preserved under small perturbations of the operator—thus any operator close enough to $\bar{\partial}$ in the appropriate functional analytic context is also surjective!

Exercise 2.47. Show that if $A : X \rightarrow Y$ is a bounded linear map between Banach spaces and $B : Y \rightarrow X$ is a bounded right inverse of A , then any small perturbation of A in the norm topology also has a bounded right inverse. *Hint: recall that any small perturbation of the identity on a Banach space is invertible, as its inverse can be expressed as a power series.*

Remark 2.48. In most presentations (e.g. [MS04], [HZ94]), some version of Prop. 2.42 and Theorem 2.43 is proven by “reducing the local problem to a global problem” so that one can apply the Fredholm theory of the Cauchy-Riemann operator. For instance, [MS04] uses the fact that $\bar{\partial}$ is a surjective Fredholm operator from $W^{1,p}$ to L^p on the *closed* unit disk if suitable boundary conditions are imposed, and a related approach is taken in [HZ94, Appendix A.4], which introduces the fundamental solution $K(z)$ and defines Tf as a convolution, but then compactifies \mathbb{C} to a sphere in order to make $\ker(\bar{\partial})$ finite dimensional. We have chosen instead to view the existence of a right inverse as an aspect of the basic local regularity theory for $\bar{\partial}$, which is a *prerequisite* for the Fredholm theory mentioned above. However, a second proof of these results will easily present itself when we discuss the Fredholm theory in Chapter 3.

2.6. Local existence of holomorphic sections

We now prove a generalization of Lemma 2.21, which implies the existence of holomorphic structures on complex vector bundles with Cauchy-Riemann type operators. The question is a purely local one, thus we can work in the trivial bundle over the open unit ball $B \subset \mathbb{C}$ with coordinates $s + it \in B$ and consider operators of the form

$$C^\infty(B, \mathbb{C}^n) \rightarrow C^\infty(B, \mathbb{C}^n) : u \mapsto \bar{\partial}u + Au$$

where $\bar{\partial}$ denotes the differential operator $\partial_s + i\partial_t$, and $A : B \rightarrow \text{End}_{\mathbb{R}}(\mathbb{C}^n)$ is a family of real linear maps on \mathbb{C}^n . For Lemma 2.21 it suffices to assume A is smooth, but in the proof and in further applications we'll find it convenient to assume that A has much weaker regularity. The smoothness of our solutions will then follow from elliptic regularity.

Theorem 2.49. *Assume $A \in L^p(B, \text{End}_{\mathbb{R}}(\mathbb{C}^n))$ for some $p \in (2, \infty]$. Then for each finite $q \in (2, p]$, there is an $\epsilon > 0$ such that for any $u_0 \in \mathbb{C}^n$, the problem*

$$\begin{aligned}\bar{\partial}u + Au &= 0 \\ u(0) &= u_0\end{aligned}$$

has a solution $u \in W^{1,q}(B_\epsilon, \mathbb{C}^n)$.

PROOF. The main idea is that if we take $\epsilon > 0$ sufficiently small, then the restriction of $\bar{\partial} + A$ to B_ϵ can be regarded as a small perturbation of the standard operator $\bar{\partial}$, and we conclude from Prop. 2.42 and Exercise 2.47 that the perturbed operator is surjective.

Since $q > 2$, the Sobolev embedding theorem implies that functions $u \in W^{1,q}$ are also continuous and bounded by $\|u\|_{W^{1,q}}$, thus we can define a bounded linear operator

$$\Phi : W^{1,q}(B) \rightarrow L^q(B) \times \mathbb{C}^n : u \mapsto (\bar{\partial}u, u(0)).$$

Prop. 2.42 implies that this operator is also surjective and has a bounded right inverse, namely

$$L^q(B) \times \mathbb{C}^n \rightarrow W^{1,q}(B) : (f, u_0) \mapsto Tf - Tf(0) + u_0,$$

where $T : L^q(B) \rightarrow W^{1,q}(B)$ is a right inverse of $\bar{\partial}$. Thus any operator sufficiently close to Φ in the norm topology also has a right inverse. Now define $\chi_\epsilon : B \rightarrow \mathbb{R}$ to be the function that equals 1 on B_ϵ and 0 outside of it, and let

$$\Phi_\epsilon : W^{1,q}(B) \rightarrow L^q(B) \times \mathbb{C}^n : u \mapsto ((\bar{\partial} + \chi_\epsilon A)u, u(0)).$$

To see that this is a bounded operator, it suffices to check that $W^{1,q} \rightarrow L^q : u \mapsto Au$ is bounded if $A \in L^p$; indeed,

$$\|Au\|_{L^q} \leq \|A\|_{L^q} \|u\|_{C^0} \leq c \|A\|_{L^p} \|u\|_{W^{1,q}},$$

again using the Sobolev embedding theorem and the assumption that $q \leq p$. Now by this same trick, we find

$$\|\Phi_\epsilon u - \Phi u\| = \|\chi_\epsilon Au\|_{L^q(B)} \leq c \|A\|_{L^p(B_\epsilon)} \|u\|_{W^{1,q}(B)},$$

thus $\|\Phi_\epsilon - \Phi\|$ is small if ϵ is small, and it follows that in this case Φ_ϵ is surjective. Our desired solution is therefore the restriction of any $u \in \Phi_\epsilon^{-1}(0, u_0)$ to B_ϵ . \square

By Exercise 2.40, the local solutions found above are smooth if $A : B \rightarrow \text{End}_{\mathbb{R}}(\mathbb{C}^n)$ is smooth, thus applying this to any smooth complex vector bundle with a complex linear Cauchy-Riemann operator, we've completed the proof of Lemma 2.21 and hence Theorem 2.20.

2.7. Similarity and unique continuation

Another consequence of the local existence result in §2.6 is that all solutions to equations of the form $\bar{\partial}u + Au = 0$, even when A is *real* linear, behave like holomorphic sections in certain respects. This will be extremely useful in studying the local properties of J -holomorphic curves, as well as global transversality issues. In practice, we'll usually need this result only in the case where A is smooth, but we'll state it in greater generality since the proof is not any harder.

Theorem 2.50 (The similarity principle). *Suppose $A \in L^\infty(B, \text{End}_{\mathbb{R}}(\mathbb{C}^n))$ and $u \in W^{1,p}(B, \mathbb{C}^n)$ for some $p > 2$ is a solution of the equation $\bar{\partial}u + Au = 0$ with $u(0) = 0$. Then for sufficiently small $\epsilon > 0$, there exist maps $\Phi \in C^0(B_\epsilon, \text{End}_{\mathbb{C}}(\mathbb{C}^n))$ and $f \in C^\infty(B_\epsilon, \mathbb{C}^n)$ such that*

$$u(z) = \Phi(z)f(z), \quad \bar{\partial}f = 0, \quad \text{and} \quad \Phi(0) = \mathbb{1}.$$

The theorem says in effect that the trivial complex vector bundle $B \times \mathbb{C}^n \rightarrow B$ admits a holomorphic structure for which the given u is a holomorphic section. In particular, this implies that if u is not identically zero, then the zero at 0 is isolated, a fact that we'll often find quite useful. There's a subtlety here to be aware of: the holomorphic structure in question is generally *not compatible with the canonical smooth structure* of the bundle, i.e. the sections that we now call “holomorphic” are not smooth in the usual sense. They will instead be of class $W^{1,q}$ for some $q > 2$, which implies they're continuous, and that's enough to imply the above statement about u having isolated zeroes. Of course, a holomorphic structure also *induces* a smooth structure on the bundle, but it will in general be a *different* smooth structure.

PROOF OF THEOREM 2.50. Given the solution $u \in W^{1,p}(B, \mathbb{C}^n)$, we claim that there exists a map $C \in L^\infty(B, \text{End}_{\mathbb{C}}(\mathbb{C}^n))$ such that $C(z)u(z) = A(z)u(z)$ almost everywhere. Indeed, whenever $u(z) \neq 0$ it is simple enough to define

$$C(z) \frac{u(z)}{|u(z)|} = A(z) \frac{u(z)}{|u(z)|}$$

and extend $C(z)$ to a complex linear map so that it satisfies a uniform bound in z almost everywhere; it need not be continuous. Now $(\bar{\partial} + C)u = 0$, and we use Theorem 2.49 to find for any $q \in (2, \infty)$ a basis of $W^{1,q}$ -smooth solutions to $(\bar{\partial} + C)v = 0$ on B_ϵ that define the standard basis of \mathbb{C}^n at 0; equivalently, this is a map $\Phi \in W^{1,q}(B_\epsilon, \text{End}_{\mathbb{C}}(\mathbb{C}^n))$ that satisfies $(\bar{\partial} + C)\Phi = 0$ and $\Phi(0) = \mathbb{1}$. Since $q > 2$, Φ is continuous and we can thus assume without loss of generality that $\Phi(z)$ is invertible everywhere on B_ϵ . Then there is a unique function $f : B_\epsilon \rightarrow \mathbb{C}^n$ such that $u = \Phi f$. A priori, f is of class $W^{1,q}$, but now the Leibnitz rule implies $\bar{\partial}f = 0$, thus f is smooth and holomorphic. \square

Exercise 2.51. By a change of local trivialization, show that a minor variation on Theorem 2.50 also holds for any $u : B \rightarrow \mathbb{C}^n$ satisfying

$$\partial_s u(z) + J(z)\partial_t u(z) + A(z)u(z) = 0,$$

where $J(z)$ is a smooth family of complex structures on \mathbb{C}^n , parametrized by $z \in B$. In particular, u has only isolated zeroes.

Remark 2.52. It will occasionally be useful to note that if the 0th-order term $A(z)$ is not only smooth but *complex* linear, then the term $\Phi(z)$ in the factorization $u(z) = \Phi(z)f(z)$ given by Theorem 2.50 will also be smooth. This is clear by a minor simplification of the proof, since it is no longer necessary to replace $A(z)$ by a separate complex linear term $C(z)$ (which in our argument above could not be assumed to be more regular than L^∞), but suffices to find a local solution of the equation $(\bar{\partial} + A)\Phi = 0$ with $\Phi(0) = \mathbb{1}$. This exists due to Theorem 2.49 and is smooth by the regularity results of §2.5. A similar remark holds in the generalized situation treated by Exercise 2.51, whenever $\partial_s + J(z)\partial_t + A(z)$ defines a *complex* linear operator with smooth coefficients, e.g. it is always true if J is smooth and $A \equiv 0$.

The following corollary will be important when we study the transversality question for global solutions to the linearized Cauchy-Riemann equation.

Corollary 2.53. *Suppose $u : (\Sigma, j) \rightarrow (M, J)$ is a smooth J -holomorphic curve and $\eta \in \Gamma(u^*TM)$ is in the kernel of the linearization $D\bar{\partial}_J(u)$. Then either $\eta \equiv 0$ or the zero set of η is discrete.*

On the local level, one can view this as a unique continuation result for J -holomorphic curves. The following is a simple special case of such a result, which we'll generalize in a moment.

Proposition 2.54. *Suppose J is a smooth almost complex structure on \mathbb{C}^n and $u, v : B \rightarrow \mathbb{C}^n$ are smooth J -holomorphic curves such that $u(0) = v(0) = 0$ and u and v have matching partial derivatives of all orders at 0. Then $u \equiv v$ on a neighborhood of 0.*

PROOF. Let $h = v - u : B \rightarrow \mathbb{C}^n$. We have

$$(2.15) \quad \partial_s u + J(u(z))\partial_t u = 0$$

and

$$(2.16) \quad \begin{aligned} \partial_s v + J(u(z))\partial_t v &= \partial_s v + J(v(z))\partial_t v + [J(u(z)) - J(v(z))]\partial_t v \\ &= -[J(u(z) + h(z)) - J(u(z))]\partial_t v \\ &= -\left(\int_0^1 \frac{d}{dt} J(u(z) + th(z)) dt\right) \partial_t v \\ &= -\left(\int_0^1 dJ(u(z) + th(z)) \cdot h(z) dt\right) \partial_t v =: -A(z)h(z), \end{aligned}$$

where the last step defines a smooth family of linear maps $A(z) \in \text{End}_{\mathbb{R}}(\mathbb{C}^n)$. Subtracting (2.15) from (2.16) gives the linear equation

$$\partial_s h(z) + J(u(z))\partial_t h(z) + A(z)h(z) = 0,$$

thus by Theorem 2.50 and Exercise 2.51, $h(z) = \Phi(z)f(z)$ near 0 for some continuous $\Phi(z) \in \text{GL}(2n, \mathbb{R})$ and holomorphic $f(z) \in \mathbb{C}^n$. Now if h has vanishing derivatives

of all orders at 0, Taylor's formula implies

$$\lim_{z \rightarrow 0} \frac{|\Phi(z)f(z)|}{|z|^k} = 0$$

for all $k \in \mathbb{N}$, so f must also have a zero of infinite order and thus $f \equiv 0$. \square

The preceding proposition is not generally as useful as one would hope, because we'll usually want to think of pseudoholomorphic curves not as specific maps but as equivalence classes of maps up to parametrization, whereas the condition that u and v have matching derivatives of all orders at a point depends heavily on the choices of parametrizations. We shall now prove a more powerful version of unique continuation that doesn't have this drawback. It will be of use to us when we study local intersection properties in §2.11.

Theorem 2.55. *Suppose j_1 and j_2 are smooth complex structures on B , J is a smooth almost complex structure on \mathbb{C}^n , and $u : (B, j_1) \rightarrow (\mathbb{C}^n, J)$ and $v : (B, j_2) \rightarrow (\mathbb{C}^n, J)$ are smooth nonconstant pseudoholomorphic curves which satisfy $u(0) = v(0) = 0$ and have matching partial derivatives to all orders at $z = 0$. Then for sufficiently small $\epsilon > 0$ there exists an embedding $\varphi : B_\epsilon \rightarrow B$ with $\varphi(0) = 0$ such that $u \equiv v \circ \varphi$ on B_ϵ .*

A corollary is that if $u, v : (B, i) \rightarrow (\mathbb{C}^n, J)$ are J -holomorphic curves that have the same ∞ -jet at 0 after a smooth reparametrization, then they are also identical up to parametrization. The reparametrization may be smooth but not necessarily holomorphic, in which case it changes i on the domain to a nonstandard complex structure j , so that the reparametrized curve no longer satisfies $\partial_s u + J(u)\partial_t u = 0$ and Prop. 2.54 thus no longer applies. We will show however that in this situation, one can find a diffeomorphism on the domain that not only transforms j back into i but also has vanishing derivatives of all orders at 0, thus producing the conditions for Prop. 2.54.

To prepare for the next lemma, recall that if $d \in \mathbb{N}$ and $u : B \rightarrow \mathbb{C}^n$ is a C^d -smooth map, then its degree d Taylor polynomial at $z = 0$ can be expressed in terms of the variables $z = s + it$ and $\bar{z} = s - it$ as

$$(2.17) \quad \sum_{k=0}^d \sum_{j+\ell=k} \frac{1}{j!\ell!} \partial_z^j \partial_{\bar{z}}^\ell u(0) z^j \bar{z}^\ell,$$

where the differential operators

$$\partial_z = \frac{\partial}{\partial z} = \frac{1}{2}(\partial_s - i\partial_t) \quad \text{and} \quad \partial_{\bar{z}} = \frac{\partial}{\partial \bar{z}} = \frac{1}{2}(\partial_s + i\partial_t)$$

are defined via the formal chain rule. If you've never seen this before, you should take a moment to convince yourself that (2.17) matches the standard Taylor's formula for a complex valued function of two real variables $u(s, t) = u(z)$. The advantage of this formalism is that it is quite easy to recognize whether a polynomial expressed in z and \bar{z} is holomorphic: the holomorphic polynomials are precisely those which only depend on powers of z , and not \bar{z} .

In the following, we'll use multiindices of the form $\alpha = (j, k)$ to denote higher order partial derivatives with respect to z and \bar{z} respectively, i.e.

$$D^\alpha = \partial_z^j \partial_{\bar{z}}^k.$$

Lemma 2.56. *Suppose $u : B \rightarrow \mathbb{C}^n$ is a smooth solution to the linear Cauchy-Riemann type equation*

$$(2.18) \quad \partial_s u(z) + J(z) \partial_t u(z) + A(z) u(z) = 0$$

with $u(0) = 0$, where $J, A \in C^\infty(B, \text{End}_{\mathbb{R}}(\mathbb{C}^n))$ with $[J(z)]^2 = -\mathbf{1}$ and $J(0) = i$. If there exists $k \in \mathbb{N}$ such that $\partial_z^\ell u(0) = 0$ for all $\ell = 1, \dots, k$, then $\partial_{\bar{z}} \partial^\alpha u(0) = 0$ for all multiindices α with $|\alpha| \leq k$. In particular, the first k derivatives of u at $z = 0$ all vanish, and $\partial_z^{k+1} u(0)$ is the only potentially nonvanishing partial derivative of order $k + 1$.

PROOF. Since $J(0) = i$, (2.18) gives $\partial_{\bar{z}} u(0) = 0$, thus we argue by induction and assume $\partial_{\bar{z}} D^\alpha u(0) = 0$ for all multiindices α of order up to $\ell \leq k - 1$. This implies that the first $\ell + 1$ derivatives of u vanish at $z = 0$. Now for any multiindex α of order $\ell + 1$, applying D^α to both sides of (2.18) and reordering the partial derivatives yields

$$\partial_s D^\alpha u(z) + J(z) \partial_t D^\alpha u(z) + \sum_{|\beta| \leq \ell+1} C_\beta(z) D^\beta u(z),$$

where $C_\beta(z)$ are smooth functions that depend on the derivatives of A and J . Evaluating at $z = 0$, the term $D^\beta u(0)$ always vanishes since $|\beta| \leq \ell + 1$, so we obtain $\partial_{\bar{z}} D^\alpha u(0) = 0$ as claimed. \square

Lemma 2.57. *Given the assumptions of Theorem 2.55, the complex structures j_1 and j_2 satisfy $j_1(0) = j_2(0)$ and also have matching partial derivatives to all orders at $z = 0$.*

PROOF. This would be obvious if u and v were immersed at 0, since then we could write $j_1 = u^* J$ and $j_2 = v^* J$, so the complex structures and their derivatives at $z = 0$ are fully determined by those of u , v and J . In general we cannot assume u and v are immersed, but we shall still use this kind of argument by taking advantage of the fact that if u and v are not constant, then Prop. 2.54 implies that they must indeed have a nonvanishing derivative of *some* order at 0.

Write $j := j_2$ and assume without loss of generality that $j_1 = i$, so u satisfies $\partial_s u + J(u) \partial_t u = 0$. We can also assume $J(0) = i$. Regarding the first derivative of u as the smooth map $du : B \rightarrow \text{Hom}_{\mathbb{R}}(\mathbb{R}^2, \mathbb{C}^n)$ defined by the matrix valued function

$$du(z) = \begin{pmatrix} \partial_s u(z) & \partial_t u(z) \end{pmatrix},$$

let $m \in \mathbb{N}$ denote the smallest order for which the m th derivative of u at $z = 0$ does not vanish. Since u also satisfies a linear Cauchy-Riemann type equation $\partial_s u + \bar{J}(z) \partial_t u$ with $\bar{J}(z) := J(u(z))$, Lemma 2.56 then implies that $\partial_z^m u(0)$ is the only nonvanishing m th order partial derivative with respect to z and \bar{z} . In particular, $\partial_z^{m-1} du(0)$ is then the lowest order nonvanishing derivative of du at $z = 0$, and the only one of order $m - 1$. We claim that the matrix $\partial_z^{m-1} du(0) \in \text{Hom}_{\mathbb{R}}(\mathbb{R}^2, \mathbb{C})$ is not

only nonzero but also nonsingular, i.e. it defines an injective linear transformation. Indeed, computing another m th order derivative of u which must necessarily vanish,

$$0 = \bar{\partial} \partial_z^{m-1} u(0) = \partial_z^{m-1} \partial_s u + i \partial_z^{m-1} \partial_t u,$$

which means that the transformation defined by $\partial_z^{m-1} du(0)$ is in fact complex linear, implying the claim.

Let us now consider together the equations satisfied by u and v :

$$\begin{aligned} du(z)i &= J(u(z))du(z), \\ dv(z)j(z) &= J(v(z))dv(z), \end{aligned}$$

where the two sides of each equation are both regarded as smooth functions $B \rightarrow \text{Hom}_{\mathbb{R}}(\mathbb{R}^2, \mathbb{C}^n)$. By assumption, the right hand sides of both equations have matching partial derivatives of all orders at $z = 0$, thus so do the left hand sides. Subtracting the second from the first, we obtain the function

$$du(z)i - dv(z)j(z) = du(z)[i - j(z)] + [du(z) - dv(z)]j(z),$$

which must have vanishing derivatives of all orders at $z = 0$. For the second term in the expression this is already obvious, so we deduce

$$(2.19) \quad D^\alpha [du \cdot (i - j)]|_{z=0} = 0$$

for all multiindices α . Applying ∂_z^{m-1} in particular and using the fact that $D^\beta du(0)$ vanishes whenever $|\beta| < m - 1$, this implies

$$\partial_z^{m-1} du(0) \cdot [i - j(0)] = 0,$$

so $j(0) = i$ since $\partial_z^{m-1} du(0)$ is injective. We now argue inductively that all higher derivatives of $i - j(z)$ must also vanish at $z = 0$. Assuming it's true for all derivatives up to order $k-1$, suppose α is a multiindex of order k , and plug the operator $\partial_z^{m-1} D^\alpha$ into (2.19). This yields

$$\partial_z^{m-1} D^\alpha [du \cdot (i - j)]|_{z=0} = c \cdot \partial_z^{m-1} du(0) \cdot D^\alpha (i - j)|_{z=0} = 0,$$

where $c > 0$ is a combinatorial constant; all other ways of distributing the operator $\partial_z^{m-1} D^\alpha$ across this product kill at least one of the two terms. Thus using the injectivity of $\partial_z^{m-1} du(0)$ once more, $D^\alpha (i - j)|_{z=0} = 0$. \square

Lemma 2.58. *Suppose j is a smooth complex structure on \mathbb{C} such that $j(0) = i$ and the derivatives $D^\alpha j(0)$ vanish for all orders $|\alpha| \geq 1$. If $\varphi : (B_\epsilon, i) \rightarrow (\mathbb{C}, j)$ is pseudoholomorphic with $\varphi(0) = 0$, then the Taylor series of φ about $z = 0$ converges to a holomorphic function on B_ϵ .*

PROOF. The map $\varphi : B_\epsilon \rightarrow \mathbb{C}$ satisfies the linear Cauchy-Riemann equation

$$(2.20) \quad \partial_s \varphi(z) + \bar{j}(z) \partial_t \varphi(z) = 0,$$

where we define $\bar{j}(z) = j(\varphi(z))$. Our conditions on j imply that $\bar{j}(0) = i$ and \bar{j} also has vanishing derivatives of all orders at 0, thus for any multiindex α , applying the differential operator D^α to both sides of (2.20) and evaluating at $z = 0$ yields $\bar{\partial} D^\alpha \varphi(0) = 0$. This implies that all terms in the Taylor expansion of φ about $z = 0$ are holomorphic, as the only nonvanishing partial derivatives in (2.17) are of the form $\partial_z^k \varphi(0)$ for $k \geq 0$.

To see that this Taylor series is actually convergent, we can use a Cauchy integral to construct the holomorphic function to which it converges: for any $\delta < \epsilon$ and $z \in B_\delta$, let

$$f(z) = \frac{1}{2\pi i} \int_{\partial B_\delta} \frac{\varphi(\zeta) d\zeta}{\zeta - z}.$$

This is manifestly a holomorphic function, and its derivatives at $z = 0$ are given by

$$(2.21) \quad f^{(n)}(0) = \frac{n!}{2\pi i} \int_{\partial B_\delta} \frac{\varphi(\zeta) d\zeta}{\zeta^{n+1}}.$$

Observe that this integral doesn't depend on the value of δ . To compute it, write φ in terms of its degree n Taylor polynomial as

$$\varphi(z) = \sum_{k=0}^n \frac{1}{k!} \partial_z^k \varphi(0) z^k + |z|^{n+1} B(z),$$

with $B(z)$ a bounded function. The integral in (2.21) thus expands into a sum of $n + 2$ terms, of which the first n are integrals of holomorphic functions and thus vanish, the last vanishes in the limit $\delta \rightarrow 0$, and the only one left is

$$f^{(n)}(0) = \frac{n!}{2\pi i} \int_{\partial B_\delta} \frac{\partial_z^n \varphi(0)}{n!} \frac{d\zeta}{\zeta} = \partial_z^n \varphi(0).$$

Thus f and φ have the same Taylor series. \square

PROOF OF THEOREM 2.55. Denote $j := j_2$ and without loss of generality, assume $j_1 \equiv i$ and $J(0) = i$. Since all complex structures on B are integrable, there exists a smooth pseudoholomorphic embedding

$$\varphi : (B, i) \rightarrow (B, j)$$

with $\varphi(0) = 0$. Now Lemma 2.57 implies that $j - i$ has vanishing derivatives of all orders at $z = 0$, and applying Lemma 2.58 in turn, we find a holomorphic function $f : B \rightarrow \mathbb{C}$ with $f(0) = 0$ whose derivatives at 0 of all orders match those of φ . In particular $f'(0) = d\varphi(0)$ is nonsingular, thus f is a biholomorphic diffeomorphism between open neighborhoods of 0, and for sufficiently small $\epsilon > 0$, we obtain a pseudoholomorphic map

$$\varphi \circ f^{-1} : (B_\epsilon, i) \rightarrow (B, j)$$

whose derivatives of all orders at 0 match those of the identity map. It follows that $v \circ \varphi \circ f^{-1} : B_\epsilon \rightarrow \mathbb{C}^n$ is now a J -holomorphic curve with the same ∞ -jet as u at $z = 0$, so Prop. 2.54 implies $v \circ \varphi \circ f^{-1} \equiv u$. \square

The similarity principle can also be used to prove certain basic facts about intersections of J -holomorphic curves: it turns out that intersection points can never accumulate except in very specific situations that are easy to characterize. We will first prove a special case of this fact, which will be generalized in §2.11 to remove the embeddedness hypothesis.

Proposition 2.59. *Suppose J is a smooth almost complex structure on \mathbb{C}^n and $u, v : B \rightarrow \mathbb{C}^n$ are smooth nonconstant J -holomorphic curves with $u(0) = v(0) = 0$. If u is also embedded, then one of the following alternatives holds:*

- For sufficiently small $\epsilon > 0$, $u(B_\epsilon \setminus \{0\}) \cap v(B_\epsilon \setminus \{0\}) = \emptyset$.
- There is a holomorphic map $\varphi : B_\epsilon \rightarrow B$ for some $\epsilon > 0$ with $\varphi(0) = 0$ and $v = u \circ \varphi$ on B_ϵ .

PROOF. By choosing coordinates intelligently, we can assume without loss of generality that $u(z) = (z, 0)$ and $J(z, 0) = i$. (Prove it!) Then writing $v(z) = (f(z), \hat{v}(z))$, one has to show that if $\hat{v} : B \rightarrow \mathbb{C}^{n-1}$ has an accumulation of zeroes near 0, it must be identically zero. By the same interpolation trick as in the proof of Prop. 2.54, we find a smooth family of linear maps $A(z) \in \text{Hom}_{\mathbb{R}}(\mathbb{C}^{n-1}, \mathbb{C}^n)$ such that

$$\partial_s v(z) + J(f(z), 0) \partial_t v(z) = \partial_s v(z) + i \partial_t v(z) = -A(z) \hat{v},$$

and thus \hat{v} satisfies $\partial_s \hat{v} + i \partial_t \hat{v} + \hat{A} \hat{v} = 0$ for some $\hat{A} \in C^\infty(B, \text{End}_{\mathbb{R}}(\mathbb{C}^{n-1}))$. The result now follows from the similarity principle. \square

2.8. Nonlinear regularity

We now extend the previous linear regularity results to the nonlinear case. In order to understand local questions regarding pseudoholomorphic maps $u : (\Sigma, j) \rightarrow (M, J)$, it suffices to study u in local coordinates near any given points on the domain and target, where by Theorem 2.6, we can always take *holomorphic* coordinates on the domain. We can therefore assume $(\Sigma, j) = (B, i)$ and M is the unit ball $B^{2n} \subset \mathbb{C}^n$, with an almost complex structure J that matches the standard complex structure i at the origin. Denote by

$$\mathcal{J}^m(B^{2n}) = \{J \in C^m(B^{2n}, \text{End}_{\mathbb{R}}(\mathbb{C}^n)) \mid J^2 \equiv -1\}$$

the space of C^m -smooth almost complex structures on B^{2n} .

Theorem 2.60. *Assume $p \in (2, \infty)$, $m \geq 1$ is an integer, $J \in \mathcal{J}^m(B^{2n})$ with $J(0) = i$ and $u : B \rightarrow B^{2n}$ is a J -holomorphic curve in $W^{1,p}(B)$ with $u(0) = 0$. Then u is also of class $W_{\text{loc}}^{m+1,p}$ on B . Moreover, if $J_k \in \mathcal{J}^m(B^{2n})$ is a sequence with $J_k \rightarrow J$ in C^m and $u_k \in W^{1,p}(B)$ is a sequence of J_k -holomorphic curves in B^{2n} converging in $W^{1,p}$ to u , then u_k also converges in $W_{\text{loc}}^{m+1,p}$.*

By the Sobolev embedding theorem, this implies that if J is smooth, then every J -holomorphic curve is also smooth, and the topology of $W_{\text{loc}}^{1,p}$ -convergence on a space of pseudoholomorphic curves is equivalent to the topology of C_{loc}^∞ -convergence. This equivalence has an important consequence for the compactness theory of holomorphic curves, arising from the fact that the hierarchy of Sobolev spaces

$$\dots \subset W^{k,p} \subset W^{k-1,p} \subset \dots \subset W^{1,p} \subset L^p$$

comes with natural inclusions that are not only continuous but also compact. Indeed, the following result plays a fundamental role in the proof of Gromov's compactness theorem, to be discussed later—it is often summarized by the phrase “gradient bounds imply C^∞ -bounds.”

Corollary 2.61. *Assume $p \in (2, \infty)$ and $m \geq 1$, $J_k \in \mathcal{J}^m(B^{2n})$ is a sequence of almost complex structures converging in C^m to $J \in \mathcal{J}^m(B^{2n})$, and $u_k : B \rightarrow B^{2n}$ is a sequence of J_k -holomorphic curves satisfying a uniform bound $\|u_k\|_{W^{1,p}(B)} < C$.*

Then u_k has a subsequence converging in $W_{\text{loc}}^{m+1,p}$ to a J -holomorphic curve $u : B \rightarrow B^{2n}$.

PROOF. Our main task is to show that u_k also satisfies a uniform bound in $W^{m+1,p}$ on every compact subset of B , as the compact embedding $W^{m+1,p} \hookrightarrow W^{m,p}$ then gives a convergent subsequence in $W_{\text{loc}}^{m,p}$, which by Theorem 2.60 must also converge in $W_{\text{loc}}^{m+1,p}$. We begin with the observation that u_k already has a C^0 -convergent subsequence, since $W^{1,p}(B)$ embeds compactly into $C^0(B)$; thus assume without loss of generality that u_k converges in C^0 to a continuous map $u : B \rightarrow B^{2n}$, and after a change of coordinates on the target, $u(0) = 0$ and $J(0) = i$.

Theorem 2.60 can be rephrased in terms of the following *local moduli spaces*: let

$$\mathcal{M}^{1,p,m} \subset C^m(B^{2n}, \text{End}_{\mathbb{R}}(\mathbb{C}^n)) \times W^{1,p}(B, \mathbb{C}^n)$$

denote the space of pairs (J, u) such that $J \in \mathcal{J}^m(B^{2n})$ and $u : B \rightarrow B^{2n}$ is a J -holomorphic curve. This is naturally a metric space due to its inclusion in the Banach space above. Similarly, for any positive number $r < 1$, define the Banach space

$$W_r^{1,p}(B, \mathbb{C}^n) = \{u \in W^{1,p}(B, \mathbb{C}^n) \mid u|_{B_r} \in W^{m+1,p}(B_r)\},$$

whose norm is the sum of the norms on $W^{1,p}(B)$ and $W^{m+1,p}(B_r)$, and define the metric subspace

$$\mathcal{M}_r^{1,p,m} = \{(J, u) \in \mathcal{J}^m(B^{2n}) \times W_r^{1,p}(B, \mathbb{C}^n) \mid u(B) \subset B^{2n} \text{ and } \partial_s u + J(u)\partial_t u = 0\}.$$

Theorem 2.60 implies that the natural inclusion

$$(2.22) \quad \mathcal{M}_r^{1,p,m} \hookrightarrow \mathcal{M}^{1,p,m}$$

is a homeomorphism. Now the pairs (J_k, u_k) form a bounded sequence in $\mathcal{M}^{1,p,m}$, and we can use the following rescaling trick to replace (J_k, u_k) by a sequence that stays within a small neighborhood of $\mathcal{J}^m(B^{2n}) \times \{0\}$. For any $\epsilon > 0$ and $u \in W^{1,p}(B)$, define the map $u^\epsilon : B \rightarrow \mathbb{C}^n$ by

$$u^\epsilon(z) = u(\epsilon z).$$

We claim that for any $\delta > 0$, one can choose $\epsilon > 0$ such that $\|u_k^\epsilon\|_{W^{1,p}(B)} < \delta$ for sufficiently large k . Indeed, integrating by change of variables,

$$\begin{aligned} \|u_k^\epsilon\|_{L^p(B)}^p &= \int_B |u_k^\epsilon(z)|^p ds dt = \frac{1}{\epsilon^2} \int_{B_\epsilon} |u_k(z)|^p ds dt \leq \frac{1}{\epsilon^2} \int_{B_\epsilon} \|u_k\|_{C^0(B_\epsilon)}^p ds dt \\ &= \pi \|u_k\|_{C^0(B_\epsilon)}^p \rightarrow \pi \|u\|_{C^0(B_\epsilon)}^p, \end{aligned}$$

where the latter is small for small ϵ since $u(0) = 0$. Likewise,

$$\|Du_k^\epsilon\|_{L^p(B)}^p = \int_B |\epsilon Du_k(\epsilon z)|^p ds dt = \epsilon^{p-2} \int_{B_\epsilon} |Du_k(z)|^p ds dt \leq \epsilon^{p-2} \|Du_k\|_{L^p(B)}^p,$$

which is small due to the uniform bound on $\|u_k\|_{W^{1,p}(B)}$. Thus choosing ϵ sufficiently small, $(J_k, u_k^\epsilon) \in \mathcal{M}^{1,p,m}$ lies in an arbitrarily small ball about $(J, 0)$ for large k , and the homeomorphism (2.22) then implies that the same is true in $\mathcal{M}_r^{1,p,m}$, thus giving a uniform bound

$$\|u_k^\epsilon\|_{W^{m+1,p}(B_r)} < C.$$

Rescaling again, this implies a uniform bound on $\|u_k\|_{W^{m+1,p}(B_{\epsilon r})}$. Since this same argument can be carried out on any sufficiently small ball about an interior point in B , and any compact subset is covered by finitely many such balls, this implies the desired bound in $W_{\text{loc}}^{m+1,p}$ on B . \square

Theorem 2.60 will be proved by induction, and the hard part is the initial step: we need to show that if J is of class C^1 , then the regularity of u can be improved from $W^{1,p}$ to $W_{\text{loc}}^{2,p}$. Observe that it suffices to find a number $\epsilon > 0$ such that $u \in W^{2,p}(B_\epsilon)$ and the sequence u_k converges in $W^{2,p}(B_\epsilon)$, since any compact subset of B can be covered by finitely many such balls of arbitrarily small radius. To obtain the desired results on B_ϵ , we will use much the same argument that was used in Prop. 2.37 for the linear case: more bookkeeping is required since J is not standard, but we'll take advantage of the assumption $J(0) = i$, so that J is *nearly* standard on B_ϵ if ϵ is sufficiently small.

PROOF OF THEOREM 2.60 FOR $m = 1$. We shall use the method of difference quotients as in Prop. 2.37 to show that $u \in W^{2,p}(B_\epsilon)$ for small $\epsilon > 0$.⁷ For any $r < 1$ and $h \in \mathbb{R} \setminus \{0\}$ sufficiently small, define a function $u^h \in W^{1,p}(B_r, \mathbb{C}^n)$ by

$$u^h(s, t) = \frac{u(s + h, t) - u(s, t)}{h},$$

so u^h converges in $L^p(B_r)$ to $\partial_s u$ as $h \rightarrow 0$. Our main goal is to find constants $\epsilon > 0$ and $C > 0$ such that

$$(2.23) \quad \|u^h\|_{W^{1,p}(B_\epsilon)} < C$$

for all sufficiently small $h \neq 0$. The Banach-Alaoglu theorem then gives a sequence $h_j \rightarrow 0$ such that u^{h_j} converges weakly in $W^{1,p}(B_\epsilon)$, implying that its limit $\partial_s u$ is also in $W^{1,p}(B_\epsilon)$; since exactly the same argument works for $\partial_t u$, we will conclude $u \in W^{2,p}(B_\epsilon)$.

To prove the bound (2.23), assume at first that ϵ is any real number with $0 < \epsilon < 1/2$; its value will be further specified later. Choose a smooth cutoff function $\beta_\epsilon : B \rightarrow [0, 1]$ with support in $B_{2\epsilon}$ such that $\beta|_{B_\epsilon} \equiv 1$. It will then suffice to show that if ϵ is taken small enough, we can find a uniform bound on $\|\beta_\epsilon u^h\|_{W^{1,p}(B_{2\epsilon})}$ as $h \rightarrow 0$. The latter has compact support in $B_{2\epsilon}$, so the Calderón-Zygmund inequality gives

$$\|\beta_\epsilon u^h\|_{W^{1,p}(B_{2\epsilon})} \leq c \|\bar{\partial}(\beta_\epsilon u^h)\|_{L^p(B_{2\epsilon})}.$$

We wish to take advantage of the fact that $\bar{\partial}_J u \equiv 0$, where we abbreviate $\bar{\partial}_J := \partial_s + J(u)\partial_t$. The latter can be regarded as the standard Cauchy-Riemann operator on a trivial bundle with nonstandard complex structure $J(u(z))$, so in particular it satisfies the Leibnitz rule $\bar{\partial}_J(fv) = (\bar{\partial}_J f)v + f(\bar{\partial}_J v)$ for $f : B \rightarrow \mathbb{R}$ and $v : B \rightarrow \mathbb{C}^n$. The difference quotient also satisfies a Leibnitz rule $(fv)^h = f^h v + f v^h$. Now rewriting $\bar{\partial}(\beta_\epsilon u^h)$ in terms of $\bar{\partial}_J$, we have

$$(2.24) \quad \bar{\partial}(\beta_\epsilon u^h) = \bar{\partial}_J(\beta_\epsilon u^h) + [i - J(u)] \partial_t(\beta_\epsilon u^h),$$

⁷The difference quotient argument explained here is adapted from the proof given in [AH, Appendix 4].

where the first term can be expanded as

$$\begin{aligned}
 \bar{\partial}_J(\beta_\epsilon u^h) &= (\bar{\partial}_J \beta_\epsilon) u^h + \beta_\epsilon \bar{\partial}_J(u^h) \\
 &= (\bar{\partial} \beta_\epsilon) u^h + [J(u) - i] (\partial_t \beta_\epsilon) u^h + \beta_\epsilon (\partial_s u^h + J(u) \partial_t u^h) \\
 (2.25) \quad &= (\bar{\partial} \beta_\epsilon) u^h + [J(u) - i] (\partial_t \beta_\epsilon) u^h + \beta_\epsilon ((\bar{\partial}_J u)^h - [J(u)]^h \partial_t u) \\
 &= (\bar{\partial} \beta_\epsilon) u^h + [J(u) - i] (\partial_t \beta_\epsilon) u^h - \beta_\epsilon [J(u)]^h \partial_t u.
 \end{aligned}$$

The last term in (2.24) satisfies the bound

$$\begin{aligned}
 \|[i - J(u)] \partial_t(\beta_\epsilon u^h)\|_{L^p(B_{2\epsilon})} &\leq \|i - J(u)\|_{C^0(B_{2\epsilon})} \|\partial_t(\beta_\epsilon u^h)\|_{L^p(B_{2\epsilon})} \\
 &\leq C_1(\epsilon) \|\beta_\epsilon u^h\|_{W^{1,p}(B_{2\epsilon})},
 \end{aligned}$$

where $C_1(\epsilon) := \|i - J(u)\|_{C^0(B_{2\epsilon})}$, and the fact that $J(u(0)) = J(0) = i$ implies that $C_1(\epsilon)$ goes to zero as $\epsilon \rightarrow 0$. We can find similar bounds for every term on the right hand side of (2.25): the first two, $\|(\bar{\partial} \beta_\epsilon) u^h\|_{L^p}$ and $\|[J(u) - i](\partial_t \beta_\epsilon) u^h\|_{L^p}$ are both bounded uniformly in h since $\|u^h\|_{L^p} \rightarrow \|\partial_s u\|_{L^p}$ as $h \rightarrow 0$. For the third term, we use the fact that $J \in C^1$ to find a pointwise bound

$$\begin{aligned}
 |[J(u)]^h(s, t)| &= \frac{1}{h} |J(u(s+h, t)) - J(u(s, t))| \leq \frac{1}{h} \|J\|_{C^1} |u(s+h, t) - u(s, t)| \\
 &= \|J\|_{C^1} |u^h(s, t)|,
 \end{aligned}$$

which implies

$$\begin{aligned}
 \|\beta_\epsilon [J(u)]^h \partial_t u\|_{L^p(B_{2\epsilon})} &\leq \|\beta_\epsilon [J(u)]^h\|_{C^0(B)} \|\partial_t u\|_{L^p(B_{2\epsilon})} \\
 &\leq C \|\beta_\epsilon u^h\|_{C^0(B)} \|u\|_{W^{1,p}(B_{2\epsilon})} \\
 &\leq C_2(\epsilon) \|\beta_\epsilon u^h\|_{W^{1,p}(B)} = C_2(\epsilon) \|\beta_\epsilon u^h\|_{W^{1,p}(B_{2\epsilon})},
 \end{aligned}$$

using the continuous embedding of $W^{1,p}(B)$ into $C^0(B)$. Here $C_2(\epsilon)$ is a constant multiple of $\|u\|_{W^{1,p}(B_{2\epsilon})}$ and thus also decays to zero as $\epsilon \rightarrow 0$. Putting all of this together, we have

$$\|\beta_\epsilon u^h\|_{W^{1,p}(B_{2\epsilon})} \leq C + C_3(\epsilon) \|\beta_\epsilon u^h\|_{W^{1,p}(B_{2\epsilon})}$$

where $C_3(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$, thus taking ϵ sufficiently small, we can move the last term to the left hand side and obtain the desired bound,

$$\|\beta_\epsilon u^h\|_{W^{1,p}(B_{2\epsilon})} \leq \frac{C}{1 - C_3(\epsilon)}.$$

The statement about convergent sequences follows by a similar argument: we assume $\|u - u_k\|_{W^{1,p}(B)} \rightarrow 0$ and use Exercise 2.36 to estimate $\|u - u_k\|_{W^{2,p}(B_\epsilon)}$ via

$$\|\beta_\epsilon(u - u_k)\|_{W^{2,p}(B_{2\epsilon})} \leq c \|\bar{\partial}(\beta_\epsilon u) - \bar{\partial}(\beta_\epsilon u_k)\|_{W^{1,p}(B_{2\epsilon})}.$$

Adapting the computation of (2.24) and (2.25) using $\partial_s u + J(u) \partial_t u = \partial_s u_k + J_k(u_k) \partial_t u_k = 0$, we find

$$\begin{aligned}
 \bar{\partial}(\beta_\epsilon u) - \bar{\partial}(\beta_\epsilon u_k) &= (\bar{\partial} \beta_\epsilon)(u - u_k) \\
 &\quad + (\partial_t \beta_\epsilon)[J(u) - i](u - u_k) + (\partial_t \beta_\epsilon)[J(u) - J_k(u_k)]u_k \\
 &\quad + [J_k(u_k) - J(u)] \partial_t(\beta_\epsilon u) + [i - J_k(u_k)][\partial_t(\beta_\epsilon u) - \partial_t(\beta_\epsilon u_k)].
 \end{aligned}$$

Since $W^{1,p}$ is a Banach algebra, it is easy to see that the first three terms in this expression each decay to zero in $W^{1,p}(B_{2\epsilon})$ as $\|u - u_k\|_{W^{1,p}} \rightarrow 0$; in particular for the third term, we use the fact that $J_k \rightarrow J$ in C^1 to conclude $J_k(u_k) \rightarrow J(u)$ in $W^{1,p}$. The fourth term is bounded similarly since $\|\partial_t(\beta_\epsilon u)\|_{W^{1,p}(B_{2\epsilon})} \leq \|\beta_\epsilon u\|_{W^{2,p}(B_{2\epsilon})}$, and we've already proved above that $u \in W^{2,p}(B_r)$ for sufficiently small r . The fifth term is a bit trickier: using the definition of the $W^{1,p}$ -norm, we have

$$\begin{aligned} \|[i - J_k(u_k)][\partial_t(\beta_\epsilon u) - \partial_t(\beta_\epsilon u_k)]\|_{W^{1,p}(B_{2\epsilon})} &\leq \|[i - J_k(u_k)][\partial_t(\beta_\epsilon u) - \partial_t(\beta_\epsilon u_k)]\|_{L^p(B_{2\epsilon})} \\ &\quad + \|DJ_k(u_k) \cdot Du_k \cdot [\partial_t(\beta_\epsilon u) - \partial_t(\beta_\epsilon u_k)]\|_{L^p(B_{2\epsilon})} \\ &\quad + \|[i - J_k(u_k)][D\partial_t(\beta_\epsilon u) - D\partial_t(\beta_\epsilon u_k)]\|_{L^p(B_{2\epsilon})}. \end{aligned}$$

In this expression, the first part is bounded by $\|\beta_\epsilon u - \beta_\epsilon u_k\|_{W^{1,p}}$, and the third by

$$\|i - J_k(u_k)\|_{C^0(B_{2\epsilon})} \|\beta_\epsilon u - \beta_\epsilon u_k\|_{W^{2,p}(B_{2\epsilon})},$$

where we observe that since $J_k(u_k(0)) \rightarrow J(u(0)) = i$, $\|i - J_k(u_k)\|_{C^0(B_{2\epsilon})}$ can be made arbitrarily small for large k by choosing ϵ sufficiently small. We use the continuous embedding $W^{1,p} \hookrightarrow C^0$ in bounding the second part by

$$\begin{aligned} \|DJ_k\|_{C^0} \|Du_k\|_{L^p(B_{2\epsilon})} \|\partial_t(\beta_\epsilon u) - \partial_t(\beta_\epsilon u_k)\|_{C^0(B)} \\ \leq c \|J_k\|_{C^1} \|u_k\|_{W^{1,p}(B_{2\epsilon})} \|\partial_t(\beta_\epsilon u) - \partial_t(\beta_\epsilon u_k)\|_{W^{1,p}(B)} \\ \leq C \|u\|_{W^{1,p}(B_{2\epsilon})} \|\beta_\epsilon u - \beta_\epsilon u_k\|_{W^{2,p}(B_{2\epsilon})}, \end{aligned}$$

where we also observe that $\|u\|_{W^{1,p}(B_{2\epsilon})}$ can be made arbitrarily small by choosing ϵ small. All together we now have a bound of the form

$$\|\beta_\epsilon(u - u_k)\|_{W^{2,p}(B_{2\epsilon})} \leq F(\|u - u_k\|_{W^{1,p}}) + C(\epsilon) \|\beta_\epsilon(u - u_k)\|_{W^{2,p}(B_{2\epsilon})}$$

for sufficiently large k , where $C(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$ and $F(t) \rightarrow 0$ as $t \rightarrow 0$. Thus choosing ϵ small enough so that $C(\epsilon) < 1$, we conclude $\|\beta_\epsilon(u - u_k)\|_{W^{2,p}(B_{2\epsilon})} \rightarrow 0$. \square

To complete the proof of Theorem 2.60 by induction, we use the following simple fact: if u is J -holomorphic, then its 1-jet can also be regarded as a pseudoholomorphic map. A global version of this statement is made precise in the appendix by P. Gauduchon of [Aud94], but we will only need a local version, which is much simpler to see. If $J \in \mathcal{J}^m(B^{2n})$, we can define an almost complex structure \hat{J} of class C^{m-1} on $B \times B^{2n} \times \mathbb{C}^n$ in block form by

$$\hat{J}(z, u, X) = \begin{pmatrix} i & 0 & 0 \\ 0 & J(u) & 0 \\ A(u, X) & 0 & J(u) \end{pmatrix},$$

where $A(u, X) \in \text{Hom}_{\mathbb{R}}(\mathbb{C}, \mathbb{C}^n)$ is defined by

$$A(u, X)(x + iy) = \begin{pmatrix} DJ(u)X \cdot X & DJ(u)X \cdot J(u)X \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

Using the fact that $0 = D(J^2)(u)X = DJ(u)X \cdot J(u) + J(u) \cdot DJ(u)X$, one can easily compute that $A(u, X)i + JA(u, X) = 0$ and thus \hat{J} is indeed an almost complex structure. Moreover, if $u : B \rightarrow B^{2n}$ satisfies $\partial_s u + J(u)\partial_t u = 0$ then

$$\hat{u} : B \rightarrow B \times B^{2n} \times \mathbb{C}^n : z \mapsto (z, u(z), \partial_s u(z))$$

satisfies $\partial_s \hat{u} + \hat{J}(\hat{u}) \partial_t \hat{u} = 0$. Indeed, this statement amounts to a system of three PDEs, of which the first is trivial, the second is $\partial_s u + J(u) \partial_t u = 0$ and the third is the latter differentiated with respect to s .

Exercise 2.62. Verify all of the above.

We can now carry out the inductive step in the proof of Theorem 2.60: assume the theorem is proved for almost complex structures of class C^{m-1} . Then if $J \in \mathcal{J}^m(B^{2n})$ and $u \in W^{1,p}(B)$ is J -holomorphic, we have $u \in W_{\text{loc}}^{m,p}$, and $\partial_s u$ is \hat{J} -holomorphic for an almost complex structure \hat{J} of class C^{m-1} , implying $\partial_s u \in W_{\text{loc}}^{m,p}$ as well. Now $\partial_t u = J(u) \partial_s u$ is also in $W_{\text{loc}}^{m,p}$ since $W^{m,p}$ is a Banach algebra, hence $u \in W_{\text{loc}}^{m+1,p}$ as claimed. The statement about converging sequences follows by a similarly simple argument.

2.9. Some tools of global analysis

To understand the structure of spaces of solutions to the nonlinear Cauchy-Riemann equation, and in particular to prove local existence in the next section, we will use the generalization of the standard differential calculus for smooth maps between Banach spaces. A readable and elegant introduction to this topic may be found in the book of Lang [Lan93]; here we shall merely summarize the essential facts.

Most of the familiar properties of derivatives and differentiable functions generalize nicely to maps between arbitrary normed linear spaces X and Y , so long as both spaces are complete. The derivative of the map $f : X \rightarrow Y$ at $x \in X$ (also often called its **linearization**) is by definition a continuous linear operator

$$df(x) \in \mathcal{L}(X, Y)$$

such that for small $h \in X$,

$$f(x+h) = f(x) + df(x)h + o(\|h\|_X),$$

where $o(\|h\|_X)$ denotes an arbitrary map of the form $\eta(h) \cdot \|h\|_X$ with $\lim_{h \rightarrow 0} \eta(h) = 0$. If $df(x)$ exists for all $x \in X$, one has a map between Banach spaces $df : X \rightarrow \mathcal{L}(X, Y)$, which may have its own derivative, and one thus obtains the notions of higher order derivatives and smoothness. Proving differentiability in the infinite dimensional setting is sometimes an intricate problem, often requiring integral inequalities such as Sobolev or Hölder estimates, and it is not hard to find natural examples of maps that are everywhere continuous but nonsmooth on some dense set.

Exercise 2.63. If $S^1 = \mathbb{R}/\mathbb{Z}$, we can denote the Banach space of real valued continuous and 1-periodic functions on \mathbb{R} by $C^0(S^1)$. Show that the map $\Phi : \mathbb{R} \times C^0(S^1) \rightarrow C^0(S^1)$ defined by $\Phi(s, f)(t) = f(s+t)$ is continuous but not differentiable.

Despite these complications, having defined the derivative, one can prove infinite dimensional versions of the familiar differentiation rules, Taylor's formula and the implicit function theorem, which can become powerful tools. The proofs, in fact, are virtually the same as in the finite dimensional case, with occasional reference to

some simple tools of linear functional analysis such as the Hahn-Banach theorem. Let us state the two most important results that we will make use of.

Theorem 2.64 (Inverse function theorem). *Suppose X and Y are Banach spaces, $\mathcal{U} \subset X$ is an open subset and $f : \mathcal{U} \rightarrow Y$ is a map of class C^k for $k \geq 1$ such that for some $x_0 \in \mathcal{U}$, $df(x_0) : X \rightarrow Y$ is a continuous isomorphism. Then f maps some neighborhood \mathcal{O} of x_0 bijectively to an open neighborhood of $y_0 := f(x_0)$, and its local inverse $f^{-1} : f(\mathcal{O}) \rightarrow \mathcal{O}$ is also of class C^k , with*

$$d(f^{-1})(y_0) = [df(x_0)]^{-1}.$$

Note that while derivatives and notions of differentiability can be defined in more general normed vector spaces, the inverse function theorem really requires X and Y to be *complete*, as the proof uses Banach's fixed point theorem (i.e. the "contraction mapping principle"). The implicit function theorem follows from this, though we should emphasize that it requires an extra hypothesis that is vacuous in the finite dimensional case:

Theorem 2.65 (Implicit function theorem). *Suppose X and Y are Banach spaces, $\mathcal{U} \subset X$ is an open subset and $f : \mathcal{U} \rightarrow Y$ is a map of class C^k for $k \geq 1$ such that for some $x_0 \in \mathcal{U}$, $df(x_0) : X \rightarrow Y$ is surjective and admits a bounded right inverse. Then there exists a C^k -map*

$$\Phi_{x_0} : \mathcal{O}_{x_0} \rightarrow X,$$

which maps some open neighborhood $\mathcal{O}_{x_0} \subset \ker df(x_0)$ of 0 bijectively to an open neighborhood of x_0 in $f^{-1}(y_0)$, where $y_0 = f(x_0)$.

Note that the existence of a bounded right inverse of $df(x_0)$ is equivalent to the existence of a splitting

$$X = \ker df(x_0) \oplus V,$$

where $V \subset X$ is a closed linear subspace, so there is a bounded linear projection map $\pi_K : X \rightarrow \ker df(x_0)$. One makes use of this in the proof as follows: assume without loss of generality that $x_0 = 0$ and consider the map

$$(2.26) \quad \Psi_0 : \mathcal{U} \rightarrow Y \oplus \ker df(0) : x \mapsto (f(x), \pi_K(x)).$$

Then $d\Psi_0(0) = (df(0), \pi_K) : X \rightarrow Y \oplus \ker df(0)$ is an isomorphism, so the inverse function theorem gives a local C^k -smooth inverse Ψ_0^{-1} , and the desired parametrization of $f^{-1}(y_0)$ can be written as $\Phi_0(v) = \Psi_0^{-1}(f(0), v)$ for sufficiently small $v \in \ker df(0)$.

Of course the most elegant way to state the implicit function theorem is in terms of manifolds: a **Banach manifold** of class C^k is simply a topological space that has local charts identifying neighborhoods with open subsets of Banach spaces such that all transition maps are C^k -smooth diffeomorphisms. Then the map Φ_{x_0} in the implicit function theorem can be regarded as the inverse of a chart, defining a Banach manifold structure on a subset of $f^{-1}(y_0)$. In fact, it is not hard to see that if $x_1, x_2 \in f^{-1}(y_0)$ are two distinct points satisfying the hypotheses of the theorem, then the resulting "transition maps"

$$\Phi_{x_1}^{-1} \circ \Phi_{x_2} : \mathcal{O}_{x_2} \rightarrow \mathcal{O}_{x_1}$$

are C^k -smooth diffeomorphisms. Indeed, these can be defined in terms of the Ψ -map of (2.26) via

$$\Psi_{x_1} \circ \Psi_{x_2}^{-1}(y_0, v) = (y_0, \Phi_{x_1}^{-1} \circ \Phi_{x_2}(v)),$$

where Ψ_{x_1} and Ψ_{x_2} are C^k -smooth local diffeomorphisms. Moreover, these charts identify the tangent space to $f^{-1}(y_0)$ at any $x_0 \in f^{-1}(y_0)$ with $\ker df(x_0) \subset X$. Thus we can restate the implicit function theorem as follows.

Corollary 2.66. *Suppose X and Y are Banach spaces, $\mathcal{U} \subset X$ is an open subset, $f : \mathcal{U} \rightarrow Y$ is a C^k -smooth map for $k \geq 1$ and $y \in Y$ is a regular value of f such that for every $x \in f^{-1}(y)$, $df(x)$ has a bounded right inverse. Then $f^{-1}(y)$ admits the structure of a C^k -smooth Banach submanifold of X , whose tangent space at $x \in f^{-1}(y)$ is $\ker df(x)$.*

By picking local charts, one sees that a similar statement is true if X and Y are also Banach manifolds instead of linear spaces, and one can generalize a step further to consider smooth sections of Banach space bundles. These results will become particularly useful when we deal with Fredholm maps, for which the linearization has finite dimensional kernel and thus satisfies the bounded right inverse assumption trivially whenever it is surjective. In this way one can prove that solution sets of certain PDEs are finite dimensional smooth manifolds. In contrast, we'll see an example in the next section of a solution set that is an *infinite dimensional* smooth Banach manifold.

The differential geometry of Banach manifolds in infinite dimensions is treated at length in [Lan99]. A more basic question is how to prove that certain spaces which naturally “should” be Banach manifolds actually are. This rather delicate question has been studied in substantial generality in the literature (see for example [Eel66, Pal68, Eli67]): the hard part is always to show that certain maps between Banach spaces are differentiable. The key is to consider only Banach spaces that have nice enough properties so that certain natural classes of maps are continuous, so that smoothness can then be proved by induction.

The next two lemmas are illustrative examples of the kinds of results one needs, and we'll make use of them in the next section. First a convenient piece of notation: if $\mathcal{U} \subset \mathbb{R}^m$ and $\Omega \subset \mathbb{R}^n$ are open subsets and $\mathbf{X}(\mathcal{U}, \mathbb{R}^n)$ denotes some Banach space of maps $\mathcal{U} \rightarrow \mathbb{R}^n$ that admits a continuous inclusion into $C^0(\mathcal{U}, \mathbb{R}^n)$, then denote

$$\mathbf{X}(\mathcal{U}, \Omega) = \{u \in \mathbf{X}(\mathcal{U}, \mathbb{R}^n) \mid u(\mathcal{U}) \subset \Omega\}.$$

Due to the continuous inclusion assumption, this is an open subset of $\mathbf{X}(\mathcal{U}, \mathbb{R}^n)$. We assume below for simplicity that Ω is convex, but this assumption is easy to remove at the cost of more cumbersome notation; see [Eli67, Lemma 4.1] for a much more general version.

Lemma 2.67. *Suppose $\mathcal{U} \subset \mathbb{R}^m$ denotes an open subset, and the symbol \mathbf{X} associates to any Euclidean space \mathbb{R}^N a Banach space $\mathbf{X}(\mathcal{U}, \mathbb{R}^N)$ consisting of bounded continuous maps $\mathcal{U} \rightarrow \mathbb{R}^N$ such that the following hypotheses are satisfied:*

- (C^0 -INCLUSION) *The inclusion $\mathbf{X}(\mathcal{U}, \mathbb{R}^N) \hookrightarrow C^0(\mathcal{U}, \mathbb{R}^N)$ is continuous.*

- (BANACH ALGEBRA) *The natural bilinear pairing*

$$\mathbf{X}(\mathcal{U}, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N)) \times \mathbf{X}(\mathcal{U}, \mathbb{R}^n) \rightarrow \mathbf{X}(\mathcal{U}, \mathbb{R}^N) : (A, u) \mapsto Au$$

is well defined and continuous.

- (C^k -CONTINUITY) *For some integer $k \geq 0$, if $\Omega \subset \mathbb{R}^n$ is any open set and $f \in C^k(\Omega, \mathbb{R}^N)$, the map*

$$(2.27) \quad \Phi_f : \mathbf{X}(\mathcal{U}, \Omega) \rightarrow \mathbf{X}(\mathcal{U}, \mathbb{R}^N) : u \mapsto f \circ u$$

is well defined and continuous.

If $\Omega \subset \mathbb{R}^n$ is a convex open set and $f \in C^{k+r}(\Omega, \mathbb{R}^N)$ for some $r \in \mathbb{N}$, then the map Φ_f defined in (2.27) is of class C^r and has derivative

$$(2.28) \quad d\Phi_f(u)\eta = (df \circ u)\eta.$$

Remark 2.68. In the formula (2.28) for the derivative we're implicitly using both the BANACH ALGEBRA and C^k -CONTINUITY hypotheses: the latter implies that $df \circ u$ is a map in $\mathbf{X}(\mathcal{U}, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N))$, which the former then embeds continuously into $\mathcal{L}(\mathbf{X}(\mathcal{U}, \mathbb{R}^n), \mathbf{X}(\mathcal{U}, \mathbb{R}^N))$.

PROOF OF LEMMA 2.67. We observe first that it suffices to prove differentiability and the formula (2.28), as $df \circ u$ is a continuous function of u and C^r -smoothness follows by induction. Thus assume $r = 1$ and $\eta \in \mathbf{X}(\mathcal{U}, \mathbb{R}^n)$ is small enough so that $u + \eta \in \mathbf{X}(\mathcal{U}, \Omega)$. Then

$$(2.29) \quad \begin{aligned} \Phi_f(u + \eta) &= \Phi_f(u) + [f \circ (u + \eta) - f \circ u] = \Phi_f(u) + \int_0^1 \frac{d}{dt} f \circ (u + t\eta) dt \\ &= \Phi_f(u) + \left[\int_0^1 df \circ (u + t\eta) dt \right] \eta \\ &= \Phi_f(u) + (df \circ u)\eta + [\theta_f \circ (u + \eta, u)] \eta, \end{aligned}$$

where we've defined $\theta_f : \Omega \times \Omega \rightarrow \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N)$ by

$$(2.30) \quad \theta_f(x, y) = \int_0^1 [df((1-t)y + tx) - df(y)] dt,$$

and observe that $\theta_f \in C^k$ since $f \in C^{k+1}$. It follows that θ_f defines a continuous map

$$\mathbf{X}(\mathcal{U}, \Omega \times \Omega) \rightarrow \mathbf{X}(\mathcal{U}, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N)) : (u, v) \mapsto \theta_f \circ (u, v),$$

and in particular

$$\lim_{\eta \rightarrow 0} \theta_f \circ (u + \eta, u) = \theta_f(u, u) = 0,$$

where the limit is taken in the topology of $\mathbf{X}(\mathcal{U}, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N))$. Thus (2.29) proves the stated formula for $d\Phi_f(u)$. \square

We will need something slightly more general, since we'll also want to be able to differentiate $(f, u) \mapsto f \circ u$ with respect to f .

Lemma 2.69. *Suppose \mathcal{U} , Ω and $\mathbf{X}(\mathcal{U}, \mathbb{R}^n)$ are as in Lemma 2.67, and in addition that the pairing $T(u)f := f \circ u$ defines T as a continuous map*

$$(2.31) \quad T : \mathbf{X}(\mathcal{U}, \Omega) \rightarrow \mathcal{L}(C^k(\Omega, \mathbb{R}^N), \mathbf{X}(\mathcal{U}, \mathbb{R}^N)).$$

Then for any $r \in \mathbb{N}$, the map

$$\Psi : C^{k+r}(\Omega, \mathbb{R}^N) \times \mathbf{X}(\mathcal{U}, \Omega) \rightarrow \mathbf{X}(\mathcal{U}, \mathbb{R}^N) : (f, u) \mapsto f \circ u$$

is of class C^r and has derivative

$$d\Psi(f, u)(g, \eta) = g \circ u + (df \circ u)\eta.$$

PROOF. We'll continue to write $\Phi_f = \Psi(f, \cdot)$ for each $f \in C^{k+r}(\Omega, \mathbb{R}^N)$; this is a C^r -smooth map $\mathbf{X}(\mathcal{U}, \Omega) \rightarrow \mathbf{X}(\mathcal{U}, \mathbb{R}^N)$ by Lemma 2.67. Observe that the pairing $T(u)f = f \circ u$ of (2.31) also gives a map

$$T : \mathbf{X}(\mathcal{U}, \Omega) \rightarrow \mathcal{L}(C^{k+r}(\Omega, \mathbb{R}^N), \mathbf{X}(\mathcal{U}, \mathbb{R}^N))$$

for each integer $r \geq 0$, and we claim that this is of class C^r . The claim mostly follows already from the proof of Lemma 2.67: expressing the remainder formula (2.29) in new notation gives

$$(2.32) \quad T(u + \eta)f = T(u)f + [T_1(u)df]\eta + [T_2(u + \eta, u)\theta_f]\eta,$$

where we've defined the related maps

$$T_1 : \mathbf{X}(\mathcal{U}, \Omega) \rightarrow \mathcal{L}(C^{k+r-1}(\Omega, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N)), \mathbf{X}(\mathcal{U}, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N))),$$

$$T_2 : \mathbf{X}(\mathcal{U}, \Omega \times \Omega) \rightarrow \mathcal{L}(C^{k+r-1}(\Omega \times \Omega, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N)), \mathbf{X}(\mathcal{U}, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N))).$$

Note that the correspondence defined in (2.30) gives a bounded linear map

$$C^{k+r}(\Omega, \mathbb{R}^N) \rightarrow C^{k+r-1}(\Omega \times \Omega, \mathcal{L}(\mathbb{R}^n, \mathbb{R}^N)) : f \mapsto \theta_f.$$

Now arguing by induction, we can assume T_1 and T_2 are both of class C^{r-1} . Then as a family of bounded linear operators acting on f , the pairing of $T_2(u + \eta, u)$ with θ_f goes to zero as $\eta \rightarrow 0$, and (2.32) implies

$$[dT(u)\eta]f = [T_1(u)df]\eta,$$

so dT is of class C^{r-1} , proving the claim.

Next consider the derivative of the map Ψ in the case $r = 1$. For any small $g \in C^{k+1}(\Omega, \mathbb{R}^N)$ and $\eta \in \mathbf{X}(\mathcal{U}, \Omega)$, we compute

$$\begin{aligned} \Psi(f + g, u + \eta) &= \Psi(f, u) + [T(u + \eta)(f + g) - T(u + \eta)(f)] + [\Phi_f(u + \eta) - \Phi_f(u)] \\ &= \Psi(f, u) + T(u)g + (T(u + \eta) - T(u))g + d\Phi_f(u)\eta + o(\|\eta\|) \\ &= \Psi(f, u) + g \circ u + (df \circ u)\eta + o(\|(g, \eta)\|). \end{aligned}$$

Thus Ψ is differentiable and we can write its derivative in the form $d\Psi(f, u) = T(u) + \Psi(df, u)$. The general result now follows easily by induction. \square

In the next section we'll apply this using the fact that if $B \subset \mathbb{C}$ is the open unit ball, then the space $W^{k,p}(B)$ (for $k \geq 1$ and $p > 2$) is a Banach algebra that embeds continuously into C^0 , and the pairing $(f, u) \mapsto f \circ u$ gives a continuous map

$$C^k(\Omega, \mathbb{R}^N) \times W^{k,p}(B, \Omega) \rightarrow W^{k,p}(B, \mathbb{R}^N).$$

Observe that by Lemma 2.67, the map $u \mapsto f \circ u$ on a suitable Banach space will be smooth if f is smooth. Things get a bit trickier if we also consider f to be a variable in this map: e.g. if f varies arbitrarily in C^k then the map $\Psi(f, u) = f \circ u$ also has only finitely many derivatives. This headache is avoided if f is allowed to vary only in some Banach space that embeds continuously into C^∞ , for then one can apply Lemma 2.69 for every k and conclude that Ψ is in C^r for all r . The most obvious examples of Banach spaces with continuous embeddings into C^∞ are finite dimensional, but we will also see an infinite dimensional example in Chapter 4 when we discuss transversality and Floer's " C_ϵ space".

2.10. Local existence of J -holomorphic curves

We shall now apply the machinery described in the previous section to prove a local existence result from which Theorem 2.6 on the integrability of Riemann surfaces follows as an easy corollary. As usual in studying such local questions, we will consider J -holomorphic maps from the unit ball $B \subset \mathbb{C}$ into $B^{2n} \subset \mathbb{C}^n$, with the coordinates chosen so that $J(0) = i$. Let B_r and B_r^{2n} denote the balls of radius $r > 0$ in \mathbb{C} and \mathbb{C}^n respectively.

In §2.1 we stated the result that there always exists a J -holomorphic curve tangent to any given vector at a given point. What we will actually prove is more general: if J is sufficiently smooth, then one can find local J -holomorphic curves with specified derivatives up to some fixed order at a point, not just the first derivative—moreover one can also find families of such curves that vary continuously under perturbations of J . Some caution is in order: it would be too much to hope that one could specify all partial derivatives arbitrarily, as the nonlinear Cauchy-Riemann equation implies nontrivial relations, e.g. $\partial_t u(0) = J(u(0))\partial_s u(0)$. What turns out to be possible is to specify the *holomorphic part* of the Taylor polynomial of u at $z = 0$ up to some finite order, i.e. the terms in the Taylor expansion that depend only on z and not on \bar{z} (cf. Equation (2.17)). The relevant higher order derivatives of u will thus be those of the form $\partial_z^k u(0)$. As the following simple result demonstrates, trying to specify more partial derivatives beyond these would yield an ill-posed problem.

Proposition 2.70. *Suppose J is a smooth almost complex structure on \mathbb{C}^n with $J(0) = i$, and $u, v : B \rightarrow \mathbb{C}^n$ are a pair of J -holomorphic curves with $u(0) = v(0) = 0$. If there exists $d \in \mathbb{N}$ such that*

$$\partial_z^k u(0) = \partial_z^k v(0)$$

for all $k = 0, \dots, d$, then in fact $D^\alpha u(0) = D^\alpha v(0)$ for every multiindex α with $|\alpha| \leq d$.

PROOF. Recall that when we used the similarity principle to prove unique continuation in Prop. 2.54, we did so by showing that $h := u - v : B \rightarrow \mathbb{C}^n$ satisfies a linear Cauchy-Riemann type equation of the form

$$\partial_s h + \bar{J}(z)\partial_t h + A(z)h = 0,$$

where in the present situation $\bar{J} : B \rightarrow \text{End}_{\mathbb{R}}(\mathbb{C}^n)$ is a smooth family of complex structures on \mathbb{C}^n and $A \in C^\infty(B, \text{End}_{\mathbb{R}}(\mathbb{C}^n))$. Since $\partial_z^k h(0) = 0$ for all $k = 0, \dots, d$, Lemma 2.56 now implies $D^\alpha h(0) = 0$ for all $|\alpha| \leq d$. \square

Here is the main local existence result.

Theorem 2.71. *Assume $p \in (2, \infty)$, $d \geq 1$ is an integer, $m \in \mathbb{N} \cup \{\infty\}$ with $m \geq d+1$, and $J \in \mathcal{J}^m(B^{2n})$ with $J(0) = i$. Then for sufficiently small $\epsilon > 0$, there exists a C^{m-d} -smooth map*

$$\Psi : (B_\epsilon^{2n})^{d+1} \rightarrow W^{d+1,p}(B, \mathbb{C}^n)$$

such that for each $(w_0, \dots, w_d) \in (B_\epsilon^{2n})^{d+1}$, $u := \Psi(w_0, \dots, w_d)$ is a J -holomorphic curve with

$$\partial_z^k u(0) = w_k$$

for each $k = 0, \dots, d$.

Exercise 2.72. Convince yourself that Theorem 2.71, together with elliptic regularity, implies that smooth almost complex structures on a real 2-dimensional manifold are always smoothly integrable, i.e. they admit smooth local charts whose transition maps are holomorphic. (See also Corollary 2.81.)

As with local existence of holomorphic sections, our proof of Theorem 2.71 will be based on the philosophy that in a sufficiently small neighborhood, everything can be understood as a perturbation of the standard Cauchy-Riemann equation. To make this precise, we will take a closer look at the *local moduli space* of J -holomorphic curves that was introduced in the proof of Corollary 2.61. For $p \in (2, \infty)$ and $k \geq 1$, define

$$W^{k,p}(B, B^{2n}) = \{u \in W^{k,p}(B, \mathbb{C}^n) \mid u(B) \subset B^{2n}\},$$

which is an open subset of $W^{k,p}(B, \mathbb{C}^n)$ due to the continuous embedding of $W^{k,p}$ in C^0 . The space of C^m -smooth almost complex structures on B^{2n} will again be denoted by $\mathcal{J}^m(B^{2n})$. Now for $J \in \mathcal{J}^m(B^{2n})$, $p \in (2, \infty)$ and $k \in \mathbb{N}$, we define the local moduli space

$$\mathcal{M}^{k,p}(J) = \{u \in W^{k,p}(B, B^{2n}) \mid \partial_s u + J(u)\partial_t u = 0\}.$$

Observe that $\mathcal{M}^{k,p}(J)$ always contains the trivial map $u \equiv 0$.

Proposition 2.73. *Suppose $J \in \mathcal{J}^m(B^{2n})$ with $J(0) = i$ and $m \geq k \geq 2$. Then some neighborhood of 0 in $\mathcal{M}^{k,p}(J)$ admits the structure of a C^{m-k+1} -smooth Banach submanifold of $W^{k,p}(B, \mathbb{C}^n)$, and its tangent space at 0 is*

$$T_0 \mathcal{M}^{k,p}(J) = \{\eta \in W^{k,p}(B, \mathbb{C}^n) \mid \bar{\partial} \eta = 0\}.$$

We prove this by presenting $\mathcal{M}^{k,p}(J)$ as the zero set of a differentiable map between Banach spaces—the tricky detail here is to determine exactly for which values of k , m and p the map in question is differentiable, and this is the essential reason behind the condition $m \geq d+1$ in Theorem 2.71. For any $p \in (2, \infty)$ and $k, m \in \mathbb{N}$ with $m \geq k-1$, let $J \in \mathcal{J}^m(B^{2n})$ with $J(0) = i$ and define the nonlinear map

$$\Phi_k : W^{k,p}(B, B^{2n}) \rightarrow W^{k-1,p}(B, \mathbb{C}^n) : u \mapsto \partial_s u + J(u)\partial_t u.$$

This is well defined due to the continuous Sobolev embedding $W^{k,p} \hookrightarrow C^{k-1}$: then $J \circ u$ is of class C^{k-1} and thus defines a bounded multiplication on $\partial_t u \in W^{k-1,p}$. One can similarly show that Φ_k is continuous, though we are much more interested in establishing conditions for it to at least C^1 .

Lemma 2.74. *If $m \geq k \geq 2$, then Φ_k is of class C^{m-k+1} , and its derivative at 0 is*

$$d\Phi_k(0) : W^{k,p}(B, \mathbb{C}^n) \rightarrow W^{k-1,p}(B, \mathbb{C}^n) : \eta \mapsto \bar{\partial}\eta.$$

PROOF. The formula for $d\Phi_k(0)$ will follow from Lemma 2.67 once we show that Φ_k is at least C^1 . The map $u \mapsto \partial_s u$ is continuous and linear, thus automatically smooth, so the nontrivial part is to show that the map $u \mapsto J(u)\partial_t u$ from $W^{k,p}(B, B^{2n})$ to $W^{k-1,p}(B, \mathbb{C}^n)$ is differentiable. Since $k \geq 2$, we can use the continuous inclusion of $W^{k,p}$ into $W^{k-1,p}$ and observe that

$$W^{k-1,p} \rightarrow W^{k-1,p} : u \mapsto J \circ u$$

is of class of C^{m-k+1} if $J \in C^m$, due to Lemma 2.67. Then differentiability of the map $u \mapsto J(u)\partial_t u$ follows from the fact that $W^{k-1,p}$ is a Banach algebra. \square

Now we apply the crucial ingredient from the linear regularity theory: Theorem 2.43 implies that $d\Phi_k(0) = \bar{\partial}$ is surjective and has a bounded right inverse. The implicit function theorem then gives $(\Phi_k)^{-1}(0)$ the structure of a differential Banach manifold near 0 and identifies its tangent space there with $\ker d\Phi_k(0) = \ker \bar{\partial}$, so the proof of Prop. 2.73 is complete.

PROOF OF THEOREM 2.71. Since $m \geq d+1$, a neighborhood of 0 in the local moduli space $\mathcal{M}^{d+1,p}(J)$ is a Banach manifold of class C^{m-d} , and $T_0\mathcal{M}^{d+1,p}(J) = \ker \bar{\partial} \subset W^{d+1,p}(B, \mathbb{C}^n)$. Due to the continuous inclusion of $W^{d+1,p}$ in C^d , there is a bounded linear *evaluation map*

$$\text{ev}_d : W^{d+1,p}(B, \mathbb{C}^n) \rightarrow (\mathbb{C}^n)^{d+1} : u \mapsto (u(0), \partial_z u(0), \partial_z^2 u(0), \dots, \partial_z^d u(0)),$$

which restricts to the local moduli space

$$\text{ev}_d : \mathcal{M}^{d+1,p}(J) \rightarrow (\mathbb{C}^n)^{d+1}$$

as a C^{m-d} -smooth map near 0. We shall use the inverse function theorem to show that ev_d maps a neighborhood of 0 in $\mathcal{M}^{d+1,p}(J)$ onto a neighborhood of 0 in $(\mathbb{C}^n)^{d+1}$ and admits a C^{m-d} -smooth right inverse.

To see this concretely, it will be convenient to restrict to a finite dimensional submanifold of $\mathcal{M}^{d+1,p}$. Let

$$\mathcal{P}_d \subset W^{d+1,p}(B, \mathbb{C}^n)$$

denote the complex $n(d+1)$ -dimensional vector space consisting of all holomorphic polynomials with degree at most d , regarded here as smooth maps $B \rightarrow \mathbb{C}^n$. Define also the closed subspace

$$\Theta^{d+1,p}(B, \mathbb{C}^n) = \text{im } \widehat{T} \subset W^{d+1,p}(B, \mathbb{C}^n),$$

where $\widehat{T} : W^{d,p}(B, \mathbb{C}^n) \rightarrow W^{d+1,p}(B, \mathbb{C}^n)$ is the bounded right inverse of $\bar{\partial} : W^{d+1,p}(B, \mathbb{C}^n) \rightarrow W^{d,p}(B, \mathbb{C}^n)$ provided by Theorem 2.43. Note that $\Theta^{d+1,p}(B, \mathbb{C}^n) \cap \mathcal{P}_d = \{0\}$ since everything in \mathcal{P}_d is holomorphic. Putting these together, we define the closed subspace

$$\Theta\mathcal{P}_d(B, \mathbb{C}^n) = \Theta^{d+1,p}(B, \mathbb{C}^n) \oplus \mathcal{P}_d \subset W^{d+1,p}(B, \mathbb{C}^n),$$

which contains an open subset

$$\Theta\mathcal{P}_d(B, B^{2n}) = \{u \in \Theta\mathcal{P}_d(B, \mathbb{C}^n) \mid u(B) \subset B^{2n}\}.$$

By construction, the restriction of $\bar{\partial} : W^{d+1,p}(B, \mathbb{C}^n) \rightarrow W^{d,p}(B, \mathbb{C}^n)$ to $\Theta\mathcal{P}_d(B, \mathbb{C}^n)$ is surjective and its kernel is precisely \mathcal{P}_d . Restricting similarly the nonlinear operator that was used to define $\mathcal{M}^{k,p}(J)$, we obtain a C^{m-d} -smooth map

$$\widehat{\Phi} : \Theta\mathcal{P}_d(B, B^{2n}) \rightarrow W^{d,p}(B, \mathbb{C}^n) : u \mapsto \partial_s u + J(u)\partial_t u,$$

whose derivative at 0 is surjective and has kernel \mathcal{P}_d , hence

$$\widehat{\mathcal{M}}(J) := \widehat{\Phi}^{-1}(0) \subset \mathcal{M}^{d+1,p}(J)$$

is a C^{m-d} -smooth finite dimensional manifold near 0, with $T_0\widehat{\mathcal{M}}(J) = \mathcal{P}_d$. Consider now the restriction of the evaluation map to $\widehat{\mathcal{M}}(J)$,

$$\text{ev}_d : \widehat{\mathcal{M}}(J) \rightarrow (\mathbb{C}^n)^{d+1}.$$

This map is linear on $W^{d+1,p}(B, \mathbb{C}^n)$, thus its derivative is simply

$$d\text{ev}_d(0) : \mathcal{P}_d \rightarrow (\mathbb{C}^n)^{d+1} : \eta \mapsto \text{ev}_d(\eta),$$

which is the isomorphism that uniquely associates to any holomorphic polynomial of degree d its derivatives of order 0 to d . Now by the inverse function theorem, the restriction of ev_d to $\widehat{\mathcal{M}}(J)$ can be inverted on a neighborhood of 0, giving rise to the desired C^{m-d} -smooth map Ψ . \square

Notice that one can extract from Theorem 2.71 *parametrized families* of local J -holomorphic curves. In particular, if $N \subset \mathbb{C}^n$ is a sufficiently small submanifold of \mathbb{C}^n , we can find a family of J -holomorphic disks $\{u_x\}_{x \in N}$ such that $u_x(0) = x$. These vary continuously in $W^{1,p}$, but actually if J is smooth, then the regularity theorem of §2.8 implies that they also vary continuously in C^∞ on compact subsets. This implies the following:

Corollary 2.75. *If J is a smooth almost complex structure on B^{2n} , $N \subset B^{2n}$ is a smooth submanifold passing through 0 and X is a smooth vector field along N , then for some neighborhood $\mathcal{U} \subset N$ of 0 and some $\epsilon > 0$, there exists a smooth family of J -holomorphic curves*

$$u_x : B \rightarrow \mathbb{C}^n, \quad x \in \mathcal{U}$$

such that $u_x(0) = x$ and $\partial_s u_x(0) = \epsilon X(x)$.

Remark 2.76. The standard meaning of the term “smooth family” as used in Cor. 2.75 is that the map $\mathcal{U} \times B \rightarrow \mathbb{C}^n : (x, z) \mapsto u_x(z)$ is smooth. Unfortunately, smoothness in this sense does not follow immediately from Theorem 2.71; the theorem rather provides smooth maps

$$\mathcal{U}_d \rightarrow W^{d,p}(B, B^{2n}) : x \mapsto u_x$$

for arbitrarily large integers $d \geq 2$ (since J is smooth), defined on open neighborhoods $\mathcal{U}_d \subset N$ whose sizes a priori depend on d . Of course more is true, as regularity guarantees that all of these maps are actually continuous into $C^\infty(B_r, B^{2n})$ for any $r < 1$, but one still must be careful in arguing that this implies a smooth family. Since we don’t have any specific applications for this result in mind, we’ll leave the details as an exercise. It should however be mentioned that this and related results

are occasionally used in the literature to construct special coordinates that make certain computations easier; see for example Exercise 2.77 below.

Exercise 2.77. Use Corollary 2.75 to show that near any point x_0 in a smooth almost complex manifold (M, J) , there exist smooth coordinates $(\zeta, w) \in \mathbb{C} \times \mathbb{C}^{n-1}$ in which $J(x_0) = i$ and in general $J(\zeta, w)$ takes the block form

$$J(\zeta, w) = \begin{pmatrix} i & Y(\zeta, w) \\ 0 & J'(\zeta, w) \end{pmatrix},$$

where $J'(\zeta, w)$ is a smooth family of complex structures on \mathbb{C}^{n-1} and $Y(\zeta, w)$ satisfies $iY + YJ' = 0$.

Finally, we can generalize local existence by allowing our local J -holomorphic curves to depend continuously on the choice of almost complex structure J . This is made possible by including $\mathcal{J}^m(B^{2n})$ into the domain of the nonlinear operator, as it will probably not surprise you to learn that the space of C^m -smooth almost complex structures is itself a smooth Banach manifold. For our purposes, it will suffice to consider small perturbations of the standard complex structure i .

Exercise 2.78. Let $\mathcal{J}(\mathbb{R}^{2n}) \subset \text{End}_{\mathbb{R}}(\mathbb{C}^n)$ denote the space of all complex structures on the vector space $\mathbb{R}^{2n} = \mathbb{C}^n$.

- (1) Consider the smooth map

$$\Phi : \text{GL}(2n, \mathbb{R}) \rightarrow \text{GL}(2n, \mathbb{R}) : A \mapsto AiA^{-1}.$$

Show that if $\text{GL}(n, \mathbb{C})$ is identified with the subgroup of all matrices in $\text{GL}(2n, \mathbb{R})$ that commute with i , then Φ descends to a bijection

$$\Phi : \text{GL}(2n, \mathbb{R}) / \text{GL}(n, \mathbb{C}) \rightarrow \mathcal{J}(\mathbb{R}^{2n}).$$

Conclude that $\mathcal{J}(\mathbb{R}^{2n})$ is a $2n^2$ -dimensional smooth submanifold of $\text{End}_{\mathbb{R}}(\mathbb{C}^n)$, and that its tangent space at i is

$$T_i \mathcal{J}(\mathbb{R}^{2n}) = \overline{\text{End}}_{\mathbb{C}}(\mathbb{C}^n) = \{Y \in \text{End}_{\mathbb{R}}(\mathbb{C}^n) \mid Yi + iY = 0\}.$$

- (2) For any $Y \in \text{End}_{\mathbb{R}}(\mathbb{C}^n)$ sufficiently small so that $\mathbb{1} + \frac{1}{2}iY$ is invertible, define

$$(2.33) \quad J_Y = \left(\mathbb{1} + \frac{1}{2}iY \right) i \left(\mathbb{1} + \frac{1}{2}iY \right)^{-1}.$$

Show that $Y \mapsto J_Y$ defines a smooth map from a neighborhood of 0 in $T_i \mathcal{J}(\mathbb{R}^{2n})$ to $\text{End}_{\mathbb{R}}(\mathbb{C}^n)$ which is bijective onto a neighborhood of i in $\mathcal{J}(\mathbb{R}^{2n})$.

- (3) If M is a smooth $2n$ -dimensional manifold with a smooth almost complex structure J , let $\text{Aut}_{\mathbb{R}}(TM)$ and $\text{Aut}_{\mathbb{C}}(TM)$ denote the open subsets of $\text{End}_{\mathbb{R}}(TM)$ and $\text{End}_{\mathbb{C}}(TM)$ respectively that consist of invertible linear maps. Show that every almost complex structure on M can then be identified with a unique section of the smooth fiber bundle

$$\text{Aut}_{\mathbb{R}}(TM) / \text{Aut}_{\mathbb{C}}(TM) \rightarrow M.$$

By Exercise 2.78, the space $\mathcal{J}^m(B^{2n})$ of C^m -smooth almost complex structures on B^{2n} can be identified with the space of C^m -smooth sections of the fiber bundle $\text{Aut}_{\mathbb{R}}(TB^{2n})/\text{Aut}_{\mathbb{C}}(TB^{2n})$, where we define $\text{Aut}_{\mathbb{C}}(TB^{2n})$ with respect to the standard complex structure of \mathbb{C}^n . One can use this fact and a version of Lemma 2.67 to show that $\mathcal{J}^m(B^{2n})$ is a smooth Banach submanifold of the Banach space $C^m(B^{2n}, \text{End}_{\mathbb{R}}(\mathbb{C}^n))$. We will not explicitly need this fact for now, but we will need a single chart: choose $\delta > 0$ sufficiently small so that (2.33) is a well defined embedding for all $|Y| < \delta$, and define the Banach space

$$\Upsilon^m = C^m(B, \overline{\text{End}_{\mathbb{C}}(\mathbb{C}^n)})$$

and open subset

$$\Upsilon_{\delta}^m = \{Y \in \Upsilon^m \mid \|Y\|_{C^0} < \delta\}.$$

Then (2.33) defines a smooth map

$$(2.34) \quad \Upsilon_{\delta}^m \rightarrow C^m(B^{2n}, \text{End}_{\mathbb{R}}(\mathbb{C}^n)) : Y \mapsto J_Y$$

which takes Υ_{δ}^m bijectively to a neighborhood of i in $\mathcal{J}^m(B^{2n})$.

Exercise 2.79. Verify that the map (2.34) is a smooth embedding. Lemma 2.67 should be useful.

Now for integers $k, m \geq 1$ and $p \in (1, \infty)$, consider the Banach space

$$X^{k,p,m} = C^m(B^{2n}, \text{End}_{\mathbb{R}}(\mathbb{C}^n)) \times W^{k,p}(B, \mathbb{C}^n)$$

and subset

$$\mathcal{M}^{k,p,m} = \{(J, u) \in \mathcal{J}^m(B^{2n}) \times W^{k,p}(B, B^{2n}) \mid \partial_s u + J(u)\partial_t u = 0\} \subset X^{k,p,m}.$$

We will call this the *local universal moduli space* of pseudoholomorphic curves. Observe that it always contains pairs of the form (i, u) where $u : B \rightarrow B^{2n}$ is holomorphic. Its local structure near such a point can be understood using the implicit function theorem: define the nonlinear map

$$\Phi_k^m : \Upsilon_{\delta}^m \times W^{k,p}(B, B^{2n}) \rightarrow W^{k-1,p}(B, \mathbb{C}^n) : (Y, u) \mapsto \partial_s u + J_Y(u)\partial_t u.$$

The zero set of this map can be identified with the space of all pairs $(J, u) \in \mathcal{M}^{k,p,m}$ such that J is within some C^m -small neighborhood of i , as then $J = J_Y$ for a unique $Y \in \Upsilon_{\delta}^m$ and $\Phi_k^m(Y, u) = 0$. Arguing as in Prop. 2.74 and applying Lemma 2.69, Φ_k^m is of class C^{m-k+1} whenever $m \geq k \geq 2$, and its derivative at any point of the form $(0, u)$ is simply

$$d\Phi_k^m(0, u)(Y, \eta) = \bar{\partial}\eta + Y(u)\partial_t u.$$

Since $\bar{\partial}$ is surjective and has a bounded right inverse, the same is always true of $d\Phi_k^m(0, u)$, and we conclude that any sufficiently small neighborhood of (i, u) in $\mathcal{M}^{k,p,m}$ is identified with a C^{m-k+1} -smooth Banach submanifold of $X^{k,p,m}$. Moreover, the natural projection

$$\pi : \mathcal{M}^{k,p,m} \rightarrow \mathcal{J}^m(B^{2n}) : (J, u) \mapsto J$$

is differentiable, and we claim that its derivative at (i, u) is also surjective, with a bounded right inverse. Indeed, identifying (i, u) with $(0, u) \in (\Phi_k^m)^{-1}(0)$, this map takes the form

$$d\pi(0, u)(Y, \eta) = Y,$$

where $(Y, \eta) \in \ker d\Phi_k^m(0, u)$ and thus satisfies the equation $\bar{\partial}\eta + Y(u)\partial_t u = 0$. Thus if $\widehat{T} : W^{k-1,p} \rightarrow W^{k,p}$ denotes a bounded right inverse of $\bar{\partial}$, then a bounded right inverse of $d\pi(0, u)$ is given by the map

$$\Upsilon^m \rightarrow \ker d\Phi_k^m(0, u) : Y \mapsto \left(Y, -\widehat{T}[Y(u)\partial_t u] \right).$$

With all of this in place, one can easily use an inversion trick as in the proof of Theorem 2.71 to show the following:

Theorem 2.80. *Suppose $u : B \rightarrow B^{2n}$ is holomorphic, i.e. it is i -holomorphic for the standard complex structure i . Then for any $p \in (2, \infty)$ and integers $m \geq k \geq 2$, there exists a neighborhood $\mathcal{U}_k^m \subset \mathcal{J}^m(B^{2n})$ of i and a C^{m-k+1} -smooth map*

$$\Psi : \mathcal{U}_k^m \rightarrow W^{k,p}(B, B^{2n})$$

such that $\Psi(0) = u$ and $\Psi(J)$ is J -holomorphic for each $J \in \mathcal{U}_k^m$.

We leave the proof as an exercise. The following simple consequence for Riemann surfaces will come in useful when we study compactness issues.

Corollary 2.81. *Suppose j_k is a sequence of complex structures on a surface Σ that converge in C^∞ to some complex structure j , and $\varphi : (B, i) \hookrightarrow (\Sigma, j)$ is a holomorphic embedding. Then for sufficiently large k , there exists a sequence of holomorphic embeddings*

$$\varphi_k : (B, i) \hookrightarrow (\Sigma, j_k)$$

that converge in C^∞ to φ .

2.11. A representation formula for intersections

In §2.7, we used the similarity principle to prove a special case of the theorem that intersections between distinct J -holomorphic curves are generally isolated: this is fairly easy if we assume one of the curves is immersed near the intersection, as we can then choose coordinates in which it takes the form $u(z) = (z, 0)$, so that intersections with the other curve $v(z) = (f(z), \hat{v}(z))$ are simply zeroes of the function $\hat{v}(z)$. In this section we will remove the assumption that u is immersed, and for this we will need a more precise description of the behavior of a J -holomorphic curve near a critical point. This description implies much more than just the fact that intersections are isolated: as we'll later see, distinct J -holomorphic curves in a 4-manifold always intersect each other *positively*, a fact which lies in the background of many powerful applications to 4-dimensional symplectic topology. We will save that discussion for later, however, and concentrate in this section on the consequences for J -holomorphic curves in arbitrary dimensions.

As a first step, it's important to understand that J -holomorphic curves have well defined tangent spaces at every point, even the critical points. Unless otherwise noted, throughout this section J will denote a smooth almost complex structure on \mathbb{C}^n with $J(0) = i$.

Proposition 2.82. *If $u : B \rightarrow \mathbb{C}^n$ is a nonconstant J -holomorphic curve with $u(0) = 0$, then there is a unique complex 1-dimensional subspace $T_u \subset \mathbb{C}^n$ and a number $k \in \mathbb{N}$ such that for every $z \in B \setminus \{0\}$, the limit*

$$\lim_{\epsilon \rightarrow 0^+} \frac{u(\epsilon z)}{\epsilon^k}$$

exists and is a nonzero vector in T_u .

PROOF. Since J is smooth, the regularity results of §2.8 imply that u is smooth, thus so is the family of complex structures defined by $\bar{J}(z) = J(u(z))$ for $z \in B$. Now u satisfies the complex linear Cauchy-Riemann type equation

$$\partial_s u + \bar{J}(z) \partial_t u = 0,$$

so by the similarity principle (see Exercise 2.51 and Remark 2.52), for sufficiently small $\delta > 0$ there is a smooth map $\Phi : B_\delta \rightarrow \text{End}_{\mathbb{R}}(\mathbb{C}^n)$ with $\Phi(0) = 1$, and a holomorphic map $f : B_\delta \rightarrow \mathbb{C}^n$ such that

$$u(z) = \Phi(z)f(z).$$

By assumption u is not constant, thus f is not identically zero and takes the form $f(z) = z^k g(z)$ for some $k \in \mathbb{N}$ and holomorphic map $g : B_\delta \rightarrow \mathbb{C}^n$ with $g(0) \neq 0$. Then for $z \in B \setminus \{0\}$ and small $\epsilon > 0$,

$$\frac{u(\epsilon z)}{\epsilon^k} = \frac{\Phi(\epsilon z) \epsilon^k z^k g(\epsilon z)}{\epsilon^k} \rightarrow z^k g(0) \in \mathbb{C}g(0)$$

as $\epsilon \rightarrow 0$. It follows that the limit of $u(\epsilon z)/\epsilon^\ell$ is either zero or infinity for all other positive integers $\ell \neq k$. \square

Definition 2.83. We will refer to the complex line $T_u \subset \mathbb{C}^n$ in Prop. 2.82 as the **tangent space** to u at 0, and its **critical order** is the integer $k - 1$.

Here now is the easiest case of the result that intersections of two different J -holomorphic curves must be isolated.

Exercise 2.84. Show that if $u, v : B \rightarrow \mathbb{C}^n$ are two nonconstant J -holomorphic curves with $u(0) = v(0) = 0$ but distinct tangent spaces $T_u \neq T_v$ at 0, then for sufficiently small $\epsilon > 0$, $u(B_\epsilon \setminus \{0\}) \cap v(B_\epsilon \setminus \{0\}) = \emptyset$.

To understand the case of an intersection with common tangency $T_u = T_v$, we will use the following local representation formula, which contains most of the hard work in this discussion.

Theorem 2.85. *For any nonconstant J -holomorphic curve $u : B \rightarrow \mathbb{C}^n$ with $u(0) = 0$, there exist smooth coordinate changes on both the domain and target, fixing the origin in both, so that in a neighborhood of 0, u is transformed into a pseudoholomorphic map $u : (B_\epsilon, \hat{j}) \rightarrow (\mathbb{C}^n, \hat{J})$, where \hat{j} and \hat{J} are smooth almost complex structures on B_ϵ and \mathbb{C}^n respectively with $\hat{j}(0) = i$ and $\hat{J}(0) = i$, and u satisfies the formula*

$$u(z) = (z^k, \hat{u}(z)) \in \mathbb{C} \times \mathbb{C}^{n-1},$$

where $k \in \mathbb{N}$ is one plus the critical order of u at 0, and $\hat{u} : B_\epsilon \rightarrow \mathbb{C}^{n-1}$ is a smooth map whose first k derivatives at 0 all vanish. In fact, \hat{u} is either identically zero or satisfies the formula

$$\hat{u}(z) = z^{k+\ell_u} C_u + |z|^{k+\ell_u} r_u(z)$$

for some constants $C_u \in \mathbb{C}^{n-1}$, $\ell_u \in \mathbb{N}$, and a function $r_u(z) \in \mathbb{C}^{n-1}$ which decays to zero as $z \rightarrow 0$.

Moreover, if $v : B \rightarrow \mathbb{C}^n$ is another nonconstant J -holomorphic curve with $v(0) = 0$ and the same tangent space and critical order as u at 0, then the coordinates above can be chosen on \mathbb{C}^n so that v (after a coordinate change on its domain) satisfies a similar representation formula $v(z) = (z^k, \hat{v}(z))$, with either $\hat{v} \equiv 0$ or $\hat{v}(z) = z^{k+\ell_v} C_v + |z|^{k+\ell_v} r_v(z)$, and any two curves written in this way are related to each other as follows: either $\hat{u} \equiv \hat{v}$, or

$$\hat{u}(z) - \hat{v}(z) = z^{k+\ell'} C' + |z|^{k+\ell'} r'(z),$$

for some constants $C' \in \mathbb{C}^{n-1}$, $\ell' \in \mathbb{N}$ and function $r'(z) \in \mathbb{C}^{n-1}$ with $\lim_{z \rightarrow 0} r'(z) = 0$. The inequality $\ell' \geq \min\{\ell_u, \ell_v\}$ holds and is strict if and only if $\ell_u = \ell_v$ and $C_u = C_v$.

Exercise 2.86. Prove Theorem 2.85 for the case where J is integrable. In this situation one can arrange for the coordinate changes on the domains to be holomorphic, so $\hat{j} \equiv i$.

Theorem 2.85 is a weak version of a deeper result proved by Micallef and White [MW95],⁸ which provides a similar formula in which the map \hat{u} can be taken to be a *polynomial* in z . That result is considerably harder to prove, but it's also more than is needed for our purposes, as the theorem above will suffice to understand everything we want to know about intersections of holomorphic curves. Before turning to the proof, let us discuss some of its applications—more will follow in a later chapter when we discuss the 4-dimensional case.

Theorem 2.87. Suppose $u, v : B \rightarrow \mathbb{C}^n$ are injective smooth J -holomorphic curves with $u(0) = v(0) = 0$. Then for sufficiently small $\epsilon > 0$, either $u = v \circ \varphi$ on B_ϵ for some holomorphic embedding $\varphi : B_\epsilon \rightarrow B$ with $\varphi(0) = 0$, or

$$u(B_\epsilon \setminus \{0\}) \cap v(B_\epsilon \setminus \{0\}) = \emptyset.$$

PROOF. By Exercise 2.84, the second alternative holds unless $T_u = T_v$, so assume the latter, and let k_u, k_v denote the critical orders of u and v respectively, plus one. Suppose $k_u m_u = k_v m_v = q$, where $q \in \mathbb{N}$ is the least common multiple of k_u and k_v , hence m_u and m_v are relatively prime. Then the two curves

$$u_0(z) := u(z^{m_u}), \quad v_0(z) := v(z^{m_v})$$

have the same tangent spaces and critical orders at 0. We can thus use Theorem 2.85 to change coordinates and rewrite these two curves as

$$u_0(z) = (z^q, \hat{u}_0(z)), \quad v_0(z) = (z^q, \hat{v}_0(z)).$$

⁸Our exposition of this topic is heavily influenced by the asymptotic version of Theorem 2.85, which is a more recent result due to R. Siefring [Sie08] that extends the intersection theory of closed J -holomorphic curves to the punctured case. We'll discuss this in a later chapter.

For each $\ell = 1, \dots, q-1$, define also the reparametrizations

$$u_\ell(z) = (z^q, \hat{u}_\ell(z)) = u_0(e^{2\pi i \ell/q} z), \quad v_\ell(z) = (z^q, \hat{v}_\ell(z)) = v_0(e^{2\pi i \ell/q} z).$$

Each of the differences $\hat{u}_0 - \hat{v}_\ell$ for $\ell = 0, \dots, q-1$ is either identically zero or satisfies a formula of the form $\hat{u}_0(z) - \hat{v}_\ell(z) = z^m C + |z|^m r(z)$, in which case it has no zeroes in some neighborhood of 0. If the latter is true for all $\ell = 0, \dots, q-1$, then u_0 has no intersections with v_0 near 0, as these correspond to pairs $z \in B_\epsilon$ and $\ell \in \{0, \dots, q-1\}$ for which $\hat{u}_0(z) = \hat{v}_\ell(z)$. It follows then that u and v have no intersections in a neighborhood of $u(0) = v(0) = 0$.

Suppose now that $\hat{u}_0 - \hat{v}_\ell \equiv 0$ for some $\ell \in \{0, \dots, q-1\}$, which means

$$(2.35) \quad u(z^{m_u}) = u_0(z) = v_0(e^{2\pi i \ell/q} z) = v(e^{2\pi i \ell/k_v} z^{m_v})$$

for all $z \in B_\epsilon$. We finish by proving the following claim: $m_u = m_v = 1$. Indeed, replacing z with $e^{2\pi i/m_u} z$ in (2.35), the left hand side doesn't change, so we deduce that for all $z \in B_\epsilon$,

$$v(z^{m_v}) = v(e^{2\pi i m_v/m_u} z^{m_v}).$$

Since v is injective by assumption, this implies $m_v/m_u \in \mathbb{Z}$, yet m_u and m_v are also relatively prime, so this can only be true if $m_u = 1$. Now performing the same argument again but inserting $e^{2\pi i/m_v} z$ into (2.35), we similarly deduce that $m_v = 1$. \square

The assumption of injectivity in the above theorem may seem like a serious restriction, but it is not: it turns out that on a sufficiently small neighborhood of each point in the domain, every nontrivial J -holomorphic curve is either injective or is a branched cover of an injective curve.

Theorem 2.88. *For any nonconstant smooth J -holomorphic curve $u : B \rightarrow \mathbb{C}^n$ with $u(0) = 0$, there exists an injective J -holomorphic curve $v : B \rightarrow \mathbb{C}^n$ and a holomorphic map $\varphi : B_\epsilon \rightarrow B$ for some $\epsilon > 0$, with $\varphi(0) = 0$, such that $u = v \circ \varphi$ on B_ϵ .*

Observe that if $\varphi'(0) \neq 0$ in the above statement then u must also be injective near 0; the interesting case is therefore when $\varphi'(0) = 0$, as then φ is locally a branched cover, mapping a neighborhood of the origin k -to-1 to another neighborhood of the origin for some $k \in \mathbb{N}$. It follows that $u : B_\epsilon \rightarrow \mathbb{C}^n$ is then also a k -fold branched cover onto the image of v near 0.

PROOF OF THEOREM 2.88. Using the coordinates provided by Theorem 2.85, rewrite u as a pseudoholomorphic map $(B_\epsilon, j) \rightarrow (\mathbb{C}^n, J)$ with $u(z) = (z^q, \hat{u}(z))$, and define for each $\ell = 0, \dots, q-1$,

$$u_\ell : (B_\epsilon, j_\ell) \rightarrow (\mathbb{C}^n, J) : z \mapsto (z^q, \hat{u}_\ell(z)) := u(e^{2\pi i \ell/q} z).$$

Then for $z \in B_\epsilon$, there is another point $\zeta \neq z$ with $u(\zeta) = u(z)$ if and only if $\hat{u}(z) = \hat{u}_\ell(z)$ for some $\ell \in \{1, \dots, q-1\}$. Making ϵ sufficiently small, the representation formula for $\hat{u} - \hat{u}_\ell$ implies that such points do not exist unless $\hat{u} \equiv \hat{u}_\ell$, so define

$$m = \min\{\ell \in \{1, \dots, q\} \mid \hat{u} \equiv \hat{u}_\ell\}.$$

Since $\hat{u} \equiv \hat{u}_m$ implies $\hat{u} \equiv \hat{u}_{\ell m}$ for all $\ell \in \mathbb{N}$, m must divide q , thus we can define a positive integer $k = q/m$. If $k = 1$ then u is injective near 0 and we are done. Otherwise, u now satisfies $u = u \circ \psi_\ell$ for all $\ell \in \mathbb{Z}_k$, where we define the diffeomorphisms

$$\psi_\ell : B_\epsilon \rightarrow B_\epsilon : z \mapsto e^{2\pi i \ell / k} z.$$

This makes it possible to define a continuous map

$$v : B_{\epsilon^k} \rightarrow \mathbb{C}^n : z \mapsto u(\sqrt[k]{z}),$$

which is injective if $\epsilon > 0$ is taken sufficiently small.

In order to view v as a J -holomorphic curve, we shall switch coordinates on the domain so that j becomes standard. Observe that since $u = u \circ \psi_\ell$, pulling J back to $\dot{B}_\epsilon := B_\epsilon \setminus 0$ through u implies $j = u^* J = \psi_\ell^* j$ on \dot{B}_ϵ for all $\ell \in \mathbb{Z}_k$, hence this holds also on B_ϵ by continuity. The maps ψ_ℓ therefore define a cyclic subgroup of the group of automorphisms of the Riemann surface (B_ϵ, j) . Find a simply connected \mathbb{Z}_k -invariant open neighborhood $\mathcal{U} \subset B_\epsilon$ of 0 which admits a holomorphic coordinate chart $\Phi : (\mathcal{U}, j) \hookrightarrow (\mathbb{C}, i)$. By the Riemann mapping theorem, we can assume without loss of generality that the image of this chart is B and $\Phi(0) = 0$, hence the inverse $\Psi := \Phi^{-1}$ defines a holomorphic embedding

$$\Psi : (B, i) \rightarrow (B_\epsilon, j)$$

that maps the origin to itself and has a \mathbb{Z}_k -invariant image. The maps

$$\tilde{\psi}_\ell := \Psi^{-1} \circ \psi_\ell \circ \Psi : (B, i) \rightarrow (B, i)$$

for $\ell \in \mathbb{Z}_k$ now define an injective homomorphism of \mathbb{Z}_k into the group of automorphisms of (B, i) that fix 0. The latter consists of rotations, so we deduce $\tilde{\psi}_\ell(z) = e^{2\pi i \ell / k} z$. Then the J -holomorphic curve $\tilde{u} := u \circ \Psi : B \rightarrow \mathbb{C}^n$ admits the symmetry $\tilde{u} = \tilde{u} \circ \tilde{\psi}_\ell$ for all $\ell \in \mathbb{Z}_k$, and we can thus define a new J -holomorphic curve on the punctured ball $\dot{B} := B \setminus \{0\}$ by

$$\tilde{v} : \dot{B} \rightarrow \mathbb{C}^n : z \mapsto \tilde{u}(\sqrt[k]{z}).$$

This admits a continuous extension over B with $\tilde{v}(0) = 0$, thus for all $z \in B_\epsilon$ in a sufficiently small neighborhood of 0, u now factors through a k -fold branched cover, namely

$$u(z) = \tilde{v}\left([\Phi(z)]^k\right).$$

Moreover, \tilde{v} is injective, which we can see by identifying it with the injective map $v : B_{\epsilon^k} \rightarrow \mathbb{C}^n$ as follows: consider the continuous map

$$f : B \rightarrow B_{\epsilon^k} : z \mapsto [\Psi(\sqrt[k]{z})]^k,$$

which is well defined because $\Psi(e^{2\pi i / k} z) = e^{2\pi i / k} \Psi(z)$. This is a homeomorphism and satisfies $\tilde{v} = v \circ f$, thus \tilde{v} is injective if and only if v is.

It remains only to show that the continuous map $\tilde{v} : B \rightarrow \mathbb{C}^n$ is in fact smooth and thus J -holomorphic at 0. By elliptic regularity (Theorem 2.60), it suffices to prove that $\tilde{v} \in W^{1,p}(B, \mathbb{C}^n)$ for some $p > 2$, i.e. that it has a weak derivative of class L^p which is defined almost everywhere and equals the smooth map $d\tilde{v}$ on \dot{B} . Recall that $u(z) = (z^q, \hat{u}(z))$ with $q = km$, where $\hat{u}(z) = o(|z|^q)$, thus the first $q - 1$

derivatives of u vanish at $z = 0$, and the same is therefore true for $\tilde{u} = u \circ \Psi$. It follows that there is a constant $C > 0$ such that

$$|d\tilde{u}(z)| \leq C|z|^{q-1}$$

for all $z \in B$, implying that for $z \in \dot{B}$,

$$|d\tilde{v}(z)| \leq |d\tilde{u}(\sqrt[k]{z})| \cdot \frac{1}{k}|z|^{\frac{1}{k}-1} \leq \frac{C}{k}|z|^{\frac{1}{k}(q-1)}|z|^{\frac{1}{k}-1} = \frac{C}{k}|z|^{m-1}.$$

Thus $d\tilde{v}$ is C^0 -bounded on \dot{B} , implying it has a finite L^p -norm for any $p > 2$, so the rest follows by Exercise 2.89 below. \square

Exercise 2.89. Assume u is any continuous function on B which is smooth on $\dot{B} = B \setminus \{0\}$, and its derivative du on \dot{B} satisfies $\|du\|_{L^p(\dot{B})} < \infty$. Show that $u \in W^{1,p}(B)$, and its weak derivative equals its strong derivative almost everywhere.

The above results are all local in nature, but they also imply the following important global consequence for *closed* holomorphic curves. Recall first that if Σ and Σ' are two closed, oriented and connected surfaces, then every continuous map

$$\varphi : \Sigma \rightarrow \Sigma'$$

has a **mapping degree** $\deg(\varphi) \in \mathbb{Z}$, most easily defined via the homological condition that $\deg(\varphi) = k$ if $\varphi_*[\Sigma] = k[\Sigma']$. Equivalently, $\deg(\varphi)$ can be defined as a signed count of points in the preimage $\varphi^{-1}(\zeta)$ of a generic point $\zeta \in \Sigma'$, cf. [Mil97].

Exercise 2.90. Show that if (Σ, j) and (Σ', j') are two closed connected Riemann surfaces with their natural orientations, then any holomorphic map $\varphi : (\Sigma, j) \rightarrow (\Sigma', j')$ has $\deg(\varphi) \geq 0$. Moreover,

- $\deg(\varphi) = 0$ if and only if φ is constant,
- $\deg(\varphi) = 1$ if and only if φ is **biholomorphic**, i.e. a holomorphic diffeomorphism with holomorphic inverse, and
- if $\deg(\varphi) = k \geq 2$, then φ is a **branched cover**, meaning it has at most finitely many critical points and its restriction to the punctured surface $\Sigma \setminus \text{Crit}(\varphi)$ is a k -fold covering map, while in a neighborhood of each critical point it admits coordinates in which $\varphi(z) = z^\ell$ for some $\ell \in \{2, \dots, k\}$.

Theorem 2.91. Suppose (Σ, j) is a closed connected Riemann surface, (M, J) is a smooth almost complex manifold and $u : (\Sigma, j) \rightarrow (M, J)$ is a nonconstant J -holomorphic curve. Then there exists a factorization $u = v \circ \varphi$ where

- (Σ', j') is a closed connected Riemann surface and $v : (\Sigma', j') \rightarrow (M, J)$ a J -holomorphic curve that is embedded outside a finite set of critical points and self-intersections, and
- $\varphi : (\Sigma, j) \rightarrow (\Sigma', j')$ is a holomorphic map of degree $\deg(\varphi) \geq 1$.

Moreover, v is unique up to biholomorphic reparametrization.

PROOF. Let $\text{Crit}(u) = \{z \in \Sigma \mid du(z) = 0\}$ denote the set of critical points, and define $\Delta \subset \Sigma$ to be the set of all points $z \in \Sigma$ such that there exists $\zeta \in \Sigma$ and neighborhoods $z \in \mathcal{U}_z \subset \Sigma$, $\zeta \in \mathcal{U}_\zeta \subset \Sigma$ with $u(z) = u(\zeta)$ but

$$u(\mathcal{U}_z \setminus \{z\}) \cap u(\mathcal{U}_\zeta \setminus \{\zeta\}) = \emptyset.$$

By Theorems 2.87 and 2.88, both of these sets are discrete and thus finite, and the set

$$\dot{\Sigma}' := u(\Sigma \setminus (\text{Crit}(u) \cup \Delta)) \subset M$$

is a smooth submanifold of M with J -invariant tangent spaces, and thus inherits a natural complex structure j' such that the inclusion $(\dot{\Sigma}', j') \hookrightarrow (M, J)$ is pseudoholomorphic. We shall now construct (Σ', j') as a compactification of $(\dot{\Sigma}', j')$, so that $\dot{\Sigma}'$ is obtained from Σ' by removing finitely many points. Let

$$\widehat{\Delta} = (\text{Crit}(u) \cup \Delta) / \sim$$

where two points in $\text{Crit}(u) \cup \Delta$ are defined to be equivalent whenever they have neighborhoods in Σ with identical images under u . Then for each $[z] \in \widehat{\Delta}$, Theorem 2.88 provides an injective J -holomorphic map $u_{[z]}$ from the open unit ball $B \subset \mathbb{C}$ onto the image of a neighborhood of z under u . We define (Σ', j') by

$$\Sigma' = \dot{\Sigma}' \cup_{\Phi} \left(\bigsqcup_{[z] \in \widehat{\Delta}} B \right),$$

where the gluing map Φ is the disjoint union of the maps $u_{[z]}|_{B \setminus \{0\}} : B \setminus \{0\} \rightarrow \dot{\Sigma}'$ for each $[z] \in \widehat{\Delta}$, and $j = j'$ on $\dot{\Sigma}'$ and i on B . The surface Σ' is clearly compact, and combining the maps $u_{[z]}$ with the inclusion $\dot{\Sigma}' \hookrightarrow M$ defines a pseudoholomorphic map $v : (\Sigma', j') \rightarrow (M, J)$ whose restriction to the punctured surface $\dot{\Sigma}' = \Sigma' \setminus \widehat{\Delta}$ is an embedding. Moreover, the restriction of u to $\Sigma \setminus (\text{Crit}(u) \cup \Delta)$ defines a holomorphic map to $(\dot{\Sigma}', j')$ which extends over the punctures to a holomorphic map $\varphi : (\Sigma, j) \rightarrow (\Sigma', j')$ such that $u = v \circ \varphi$.

We leave the uniqueness statement as an exercise for the reader. The positivity of $\deg(\varphi)$ follows from Exercise 2.90. \square

Finally, we turn to the proof of the representation formula, Theorem 2.85. A somewhat simplified characterization of the argument would be as follows: we need to show that for any nonconstant J -holomorphic curve $u : B \rightarrow \mathbb{C}^n$, assuming $J(0) = i$, the “leading order” terms in its Taylor expansion about $z = 0$ are holomorphic. Since terms in the Taylor series can always be expressed as constant multiples of $z^k \bar{z}^\ell$, holomorphicity means the relevant terms are actually multiples of z^k , thus producing the powers of z that appear in the representation formula. In practice, things are a bit more complicated than this, e.g. to keep full control over the remainders, we will at one point use the similarity principle instead of Taylor’s theorem, but the above can be seen as a motivating principle.

PROOF OF THEOREM 2.85. We proceed in four steps.

Step 1: Coordinates on the target. Choose the coordinates on \mathbb{C}^n so that $J(0) = i$ and $T_u = \mathbb{C} \times \{0\} \subset \mathbb{C}^n$. We can make one more requirement on the coordinates without loss of generality: we choose them so that the map

$$u_0(z) = (z, 0) \in \mathbb{C} \times \mathbb{C}^{n-1}$$

is J -holomorphic on B_ϵ for sufficiently small $\epsilon > 0$. This is a highly nontrivial condition: the fact that it's possible follows from the local existence result for J -holomorphic curves with a fixed tangent vector, Theorem 2.71.

Step 2: Coordinates on the domain. We next seek a coordinate change near the origin on the domain so that u becomes a map of the form $z \mapsto (z^k, o(|z|^k))$ for some $k \in \mathbb{N}$. Applying the similarity principle as in the proof of Prop. 2.82, we have $u(z) = \Phi(z)f(z)$ on B_ϵ for some small $\epsilon > 0$, a smooth map $\Phi : B_\epsilon \rightarrow \text{End}_{\mathbb{R}}(\mathbb{C}^n)$ with $\Phi(0) = \mathbb{1}$ and a holomorphic map $f : B_\epsilon \rightarrow \mathbb{C}^n$. Moreover, $f(z) = z^k g(z)$ for some $k \in \mathbb{N}$ (where $k-1$ is the critical order of u) and a holomorphic map $g : B_\epsilon \rightarrow \mathbb{C}^n$ with $g(0) \neq 0$, and our assumption on T_u implies that after a complex linear coordinate change on the domain, we may assume $g(0) = (1, 0) \in \mathbb{C} \times \mathbb{C}^{n-1}$. Thus $f(z) = (z^k g_1(z), z^{k+1} g_2(z))$ for some holomorphic maps $g_1 : B_\epsilon \rightarrow \mathbb{C}$ and $g_2 : B_\epsilon \rightarrow \mathbb{C}^{n-1}$, with $g_1(0) = 1$. Let us use the splitting $\mathbb{C}^n = \mathbb{C} \times \mathbb{C}^{n-1}$ to write $\Phi(z)$ in block form as

$$\Phi(z) = \begin{pmatrix} \alpha(z) & \beta(z) \\ \gamma(z) & \delta(z) \end{pmatrix},$$

so $\alpha(0)$ and $\delta(0)$ are both the identity, while $\beta(0)$ and $\gamma(0)$ both vanish; note that all four blocks are regarded as *real* linear maps on complex vector spaces, i.e. they need not commute with multiplication by i . Now $u(z)$ takes the form $(u_1(z), u_2(z)) \in \mathbb{C} \times \mathbb{C}^{n-1}$, where

$$\begin{aligned} u_1(z) &= \alpha(z)z^k g_1(z) + \beta(z)z^{k+1} g_2(z), \\ u_2(z) &= \gamma(z)z^k g_1(z) + \delta(z)z^{k+1} g_2(z). \end{aligned}$$

We claim that after shrinking $\epsilon > 0$ further if necessary, there exists a smooth function $\zeta : B_\epsilon \rightarrow \mathbb{C}$ such that $\zeta(0) = 0$, $d\zeta(0) = \mathbb{1}$ and $[\zeta(z)]^k = u_1(z)$. Indeed, the desired function can be written as

$$\zeta(z) = z \sqrt[k]{\alpha(z)g_1(z) + \beta(z)zg_2(z)},$$

which can be defined as a smooth function for z near 0 since the expression under the root lies in a neighborhood of 1; we set $\sqrt[k]{1} = 1$. Expressing u now as a function of the new coordinate ζ , we have

$$(2.36) \quad u(\zeta) = (\zeta^k, \hat{u}(\zeta))$$

with $\hat{u}(\zeta) = A(\zeta)\zeta^k$ for some smooth map $A(\zeta) \in \text{Hom}_{\mathbb{R}}(\mathbb{C}, \mathbb{C}^{n-1})$ with $A(0) = 0$. Observe that since $d\zeta(0) = \mathbb{1}$, the new expression for $u(\zeta)$ is pseudoholomorphic for a new complex structure \hat{j} on the domain such that $\hat{j}(0) = i$.

Step 3: The leading order term in $\hat{u} - \hat{v}$. This is the important part. Using the coordinates chosen above, assume now that $J(0) = i$ and the two maps $u : (B_\epsilon, j) \rightarrow (\mathbb{C}^n, J)$ and $v : (B_\epsilon, j') \rightarrow (\mathbb{C}^n, J)$ are pseudoholomorphic curves of the form

$$\begin{aligned} u(z) &= (z^k, \hat{u}(z)), \\ v(z) &= (z^k, \hat{v}(z)), \end{aligned}$$

where \hat{u} and \hat{v} each have vanishing derivatives up to at least order k at $z = 0$. Let

$$h(z) = u(z) - v(z) = (0, \hat{h}(z)),$$

defining a map $\hat{h} : B_\epsilon \rightarrow \mathbb{C}^{n-1}$. Our main goal is to show that the leading order term in \hat{h} is a homogeneous holomorphic polynomial. By unique continuation (Theorem 2.55), h vanishes identically on a neighborhood of 0 if and only if the derivatives $D^\ell h(0)$ of all orders vanish, so let's assume this is not the case. Then there is a finite positive integer m defined by

$$m = \min\{\ell \in \mathbb{N} \mid D^\ell h(0) \neq 0\},$$

and $m \geq k + 1$ since $h(z) = o(|z|^k)$. Now for $\epsilon > 0$, the functions

$$h_\epsilon(z) := \frac{h(\epsilon z)}{\epsilon^m}$$

converge in C^∞ as $\epsilon \rightarrow 0$ to a nonzero homogeneous polynomial in z and \bar{z} of degree m , namely the m th order term in the Taylor series of h about 0. We claim that this polynomial is holomorphic, which would imply that it has the form

$$h_0(z) = (0, z^m C)$$

for some constant $C \in \mathbb{C}^{n-1}$.

The intuitive reason for this claim should be clear: u and v both satisfy nonlinear Cauchy-Riemann equations that “converge” to the standard one as $z \rightarrow 0$, so their difference in the rescaled limit should also satisfy $\bar{\partial}h_0 = 0$. One complication in making this argument precise is that since we've reparametrized the domains by nonholomorphic diffeomorphisms, u and v are each pseudoholomorphic for *different complex structures* j and j' on their domains, thus it is not so straightforward to find an appropriate PDE satisfied by $u - v$. Of course, since both maps are immersed except at 0, the complex structures are uniquely determined by $j = u^*J$ and $j' = v^*J$ on $B_\epsilon \setminus \{0\}$, which suggests that there should be a way to reexpress the two nonlinear Cauchy-Riemann equations without explicit reference to j and j' . And there is: we only need observe that outside of 0, u and v parametrize immersed surfaces in \mathbb{C}^n whose tangent spaces are *complex*, i.e. J -invariant.

This can be expressed elegantly in the language of **bivectors**: recall that a bivector is an element of the antisymmetric tensor product bundle $\Lambda^2 T\mathbb{C}^n \rightarrow \mathbb{C}^n$, and thus consists of a linear combination of bilinear wedge products of the form $X \wedge Y$ for vectors $X, Y \in T_p\mathbb{C}^n$, $p \in \mathbb{C}^n$, where by definition $X \wedge Y = -Y \wedge X$. Such a product can be thought of intuitively as representing the oriented linear subspace in $T_p\mathbb{C}^n$ spanned by X and Y , with its magnitude giving the signed area of the corresponding parallelogram. Let $\text{Aut}_{\mathbb{R}}(E)$ denote the group of invertible real linear smooth bundle maps on any bundle E . Then there is a natural group homomorphism

$$\text{Aut}_{\mathbb{R}}(T\mathbb{C}^n) \rightarrow \text{Aut}_{\mathbb{R}}(\Lambda^2 T\mathbb{C}^n) : A \mapsto \bar{A}$$

defined by

$$\bar{A}(X \wedge Y) = AX \wedge AY.$$

In particular, $J^2 = -1$ then implies $\bar{J}^2 = 1$ as an operator on $\Lambda^2 T\mathbb{C}^n$. Now, the action of J fixes the oriented subspace spanned by X and Y if and only if $JX \wedge JY = c(X \wedge Y)$ for some $c > 0$, but from $\bar{J}^2 = 1$, we deduce that $c = 1$, so

the correct condition is $JX \wedge JY = X \wedge Y$. We conclude from this discussion that $u : B_\epsilon \rightarrow \mathbb{C}^n$ and $v : B_\epsilon \rightarrow \mathbb{C}^n$ satisfy the first order nonlinear PDEs,

$$(2.37) \quad \begin{aligned} \partial_s u \wedge \partial_t u - J(u) \partial_s u \wedge J(u) \partial_t u &= 0, \\ \partial_s v \wedge \partial_t v - J(v) \partial_s v \wedge J(v) \partial_t v &= 0. \end{aligned}$$

In order to deduce the consequence for h_0 , observe first that by the usual interpolation trick (cf. the proof of Prop. 2.54), on a sufficiently small ball B_ϵ there is a smooth map $A : B_\epsilon \rightarrow \text{End}_{\mathbb{R}}(\mathbb{C}^n, \text{End}_{\mathbb{R}}(\mathbb{C}^n))$ such that

$$J(u(z)) - J(v(z)) = A(z) [u(z) - v(z)] = A(z)h(z).$$

Thus subtracting the second equation of (2.37) from the first gives

$$\begin{aligned} \partial_s u \wedge \partial_t h + \partial_s h \wedge \partial_t v - J(u) \partial_s u \wedge J(u) \partial_t h - J(u) \partial_s h \wedge J(v) \partial_t v \\ - J(u) \partial_s u \wedge (Ah) \partial_t v - (Ah) \partial_s v \wedge J(v) \partial_t v = 0. \end{aligned}$$

Replacing z by ϵz and dividing the entire expression by ϵ^{k+m-2} now yields

$$\begin{aligned} 0 &= \frac{\partial_s u(\epsilon z)}{\epsilon^{k-1}} \wedge \frac{\partial_t h(\epsilon z)}{\epsilon^{m-1}} + \frac{\partial_s h(\epsilon z)}{\epsilon^{m-1}} \wedge \frac{\partial_t v(\epsilon z)}{\epsilon^{k-1}} \\ &\quad - J(u(\epsilon z)) \frac{\partial_s u(\epsilon z)}{\epsilon^{k-1}} \wedge J(u(\epsilon z)) \frac{\partial_t h(\epsilon z)}{\epsilon^{m-1}} - J(u(\epsilon z)) \frac{\partial_s h(\epsilon z)}{\epsilon^{m-1}} \wedge J(v(\epsilon z)) \frac{\partial_t v(\epsilon z)}{\epsilon^{k-1}} \\ &\quad - \epsilon^k J(u(\epsilon z)) \frac{\partial_s u(\epsilon z)}{\epsilon^{k-1}} \wedge \left[A(\epsilon z) \frac{h(\epsilon z)}{\epsilon^m} \right] \frac{\partial_t v(\epsilon z)}{\epsilon^{k-1}} \\ &\quad - \epsilon^k \left[A(\epsilon z) \frac{h(\epsilon z)}{\epsilon^m} \right] \frac{\partial_s v(\epsilon z)}{\epsilon^{k-1}} \wedge J(v(\epsilon z)) \frac{\partial_t v(\epsilon z)}{\epsilon^{k-1}}. \end{aligned}$$

We claim that every term in this expression converges in C^∞ as $\epsilon \rightarrow 0$. Indeed, the terms involving h are all either $h_\epsilon(z)$ or one of its first derivatives, so these converge respectively to $h_0 = (0, \hat{h}_0)$, $\partial_s h_0 = (0, \partial_s \hat{h}_0)$ and $\partial_t h_0 = (0, \partial_t \hat{h}_0)$. Since $\partial_s u$ has vanishing derivatives at 0 up until order $k-1$, $\frac{\partial_s u(\epsilon z)}{\epsilon^{k-1}}$ converges to the homogenous degree $k-1$ Taylor polynomial of $\partial_s u$ at 0, which is precisely the first derivative of the leading order term in u , namely $(kz^{k-1}, 0)$. Likewise, $\frac{\partial_t u(\epsilon z)}{\epsilon^{k-1}} \rightarrow (ikz^{k-1}, 0)$, and the same goes for the first derivatives of v . Finally $J(u(\epsilon z))$ and $J(v(\epsilon z))$ both converge to i , so after the dust settles, we're left with

$$\begin{aligned} (kz^{k-1}, 0) \wedge (0, \partial_t \hat{h}_0) + (0, \partial_s \hat{h}_0) \wedge (ikz^{k-1}, 0) \\ - (ikz^{k-1}, 0) \wedge (0, i\partial_t \hat{h}_0) + (0, i\partial_s \hat{h}_0) \wedge (kz^{k-1}, 0) = 0, \end{aligned}$$

or equivalently

$$-(kz^{k-1}, 0) \wedge (0, i\bar{\partial} \hat{h}_0) = (ikz^{k-1}, 0) \wedge (0, \bar{\partial} \hat{h}_0).$$

This equation means that for all $z \in B_\epsilon$, if $(kz^{k-1}, 0)$ and $(0, i\bar{\partial} h_0(z))$ are linearly independent vectors in \mathbb{C}^n , then the oriented real subspace they span is the same as its image under multiplication by i , i.e. it is complex. But this is manifestly untrue unless one of the vectors vanishes, so we conclude that for all $z \in B_\epsilon \setminus \{0\}$, $\bar{\partial} h_0(z) = 0$, and h_0 is thus a holomorphic polynomial on B_ϵ .

Step 4: Conclusion. It remains only to assemble the information gathered above. Combining Step 3 with Taylor's theorem yields the expression

$$\hat{u}(z) - \hat{v}(z) = z^m C + |z|^m r(z),$$

where $C \in \mathbb{C}^{n-1}$ is a constant, $m > k$ is an integer and $r(z)$ is a remainder function such that $\lim_{z \rightarrow 0} r(z) = 0$. The corresponding formulas for \hat{u} and \hat{v} individually follow from this, because we've chosen coordinates so that $z \mapsto u_0(z^k) = (z^k, 0)$ is also a J -holomorphic curve. The degree of the leading term in each is then simply the degree of its lowest order nonvanishing derivative at $z = 0$, and the same applies to $\hat{u} - \hat{v}$. \square

CHAPTER 3

Fredholm theory

3.1. Some Banach spaces and manifolds

In this chapter we begin the study of J -holomorphic curves in global settings. We will fix the following data throughout: (Σ, j) is a closed connected Riemann surface, and (M, J) is a $2n$ -dimensional manifold with a smooth almost complex structure. Our goal will be to understand the local structure of the space of solutions to the nonlinear Cauchy-Riemann equation, that is,

$$(3.1) \quad \{u \in C^\infty(\Sigma, M) \mid Tu \circ j = J \circ Tu\}.$$

We assign to this space the natural topology defined by C^∞ -convergence of maps $\Sigma \rightarrow M$. Recall that since J is smooth, elliptic regularity implies that all solutions of at least class $W_{\text{loc}}^{k,p}$ for some $k \in \mathbb{N}$ and $p > 2$ are actually smooth, and the C^∞ -topology is equivalent to the $W^{k,p}$ -topology on the solution space. The main result of this chapter will be that under sufficiently fortunate circumstances, this space is a finite dimensional manifold, and we will compute its dimension in terms of the given topological data. We will put off until the next chapter the question of when such “fortunate circumstances” are guaranteed to exist, i.e. when *transversality* is achieved. It should also be noted that in later applications to symplectic topology, the space (3.1) will not really be the one we want to consider: it has two unnatural features, namely that it fixes an *arbitrary* complex structure on the domain, and that it may include different curves that are reparametrizations of each other, and thus should really be considered “equivalent”. We will address these issues in Chapter 4, when we give the proper definition of the moduli space of J -holomorphic curves.

For now, (3.1) will be the space of interest, and we sketched already in §2.4 how to turn the study of this space into a problem of nonlinear functional analysis. It is time to make that discussion precise by defining the appropriate Banach manifolds and bundles.

We must first understand how to define Sobolev spaces of sections on vector bundles. In general, for any smooth vector bundle $E \rightarrow \Sigma$ one can define the space $W_{\text{loc}}^{k,p}(E)$ to consist of all sections whose expressions in all choices of local coordinates and trivializations are of class $W^{k,p}$ on compact subsets. One can analogously define maps of class $W_{\text{loc}}^{k,p}$ between two smooth manifolds. When Σ is also *compact*, we define the space $W^{k,p}(E)$ to be simply $W_{\text{loc}}^{k,p}(E)$, and give it the structure of a Banach space as follows. Choose a finite open cover $\bigcup_j \mathcal{U}_j = \Sigma$, and assume that for each set $\mathcal{U}_j \subset \Sigma$, there is a smooth chart $\varphi_j : \mathcal{U}_j \rightarrow \Omega_j$, where $\Omega_j = \varphi_j(\mathcal{U}_j) \subset \mathbb{C}$, as well as a local trivialization $\Phi_j : E|_{\mathcal{U}_j} \rightarrow \mathcal{U}_j \times \mathbb{C}^n$. Then if $\{\alpha_j : \Sigma \rightarrow [0, 1]\}$ is a

partition of unity subordinate to $\{\mathcal{U}_j\}$, define for any section $v : \Sigma \rightarrow E$,

$$(3.2) \quad \|v\|_{W^{k,p}(E)} = \sum_j \|\text{pr}_2 \circ \Phi_j \circ (\alpha_j v) \circ \varphi_j^{-1}\|_{W^{k,p}(\Omega_j)}.$$

This definition depends on plenty of choices, and the norm on $W^{k,p}(E)$ is thus not canonically defined; really one should call $W^{k,p}(E)$ a *Banachable* space rather than a Banach space. The exercise below shows that at least the resulting *topology* on $W^{k,p}(E)$ is canonical. In a completely analogous way, one can also define the Banach spaces $C^k(E)$ and $C^{k,\alpha}(E)$.

Exercise 3.1.

- (1) Show that any alternative choice of finite open covering, charts, trivializations and partition of unity gives an equivalent norm on $W^{k,p}(E)$. *Hint: given two complete norms on the same vector space, it's enough to show that the identity map from one to the other is continuous (in one direction!).*
- (2) Verify that your favorite embedding theorems hold: in particular, $W^{k,p}(E)$ admits a continuous and compact embedding into $W^{k-1,p}(E)$ for all p , and $C^{k-1}(E)$ whenever $p > 2$.

Remark 3.2. If Σ is not compact, then even the topology of $W^{k,p}(E)$ is not well defined without some extra choices. We'll need to deal with this issue later when we discuss punctured holomorphic curves.

Exercise 3.3. For $kp > 2$, a vector bundle $E \rightarrow \Sigma$ is said to have a $W^{k,p}$ -**smooth structure** if it admits a system of local trivializations whose transition maps are of class $W^{k,p}$. Show that $W^{k,p}(E)$ is also a well defined Banachable space in this case, though one cannot speak of sections of any better regularity than $W^{k,p}$. Why doesn't any of this make sense if $kp \leq 2$?

Next we consider maps of Sobolev-type regularity between the manifolds Σ and M ; we'll restrict our attention to the case $kp > 2$, so that all such maps are continuous. It was already remarked that the space $W_{\text{loc}}^{k,p}(\Sigma, M)$ can be defined naturally by expressing maps $\Sigma \rightarrow M$ in local charts, though since it isn't a vector space, the question of precisely what structure this space has is a bit subtle. Intuitively, we expect spaces of maps $\Sigma \rightarrow M$ to be manifolds, and this motivates the following definition.

Definition 3.4. For any $k \in \mathbb{N}$ and $p > 1$ such that $kp > 2$, choose any smooth connection on M , and for any smooth map $f \in C^\infty(\Sigma, M)$, choose a neighborhood \mathcal{U}_f of the zero section in f^*TM such that for all $z \in \Sigma$, the restriction of \exp to $T_{f(z)}M \cap \mathcal{U}_f$ is an embedding. Then we define the space of $W^{k,p}$ -smooth maps from Σ to M by

$$W^{k,p}(\Sigma, M) = \{u \in C^0(\Sigma, M) \mid u = \exp_f \eta \text{ for any } f \in C^\infty(\Sigma, M) \text{ and } \eta \in W^{k,p}(f^*TM) \text{ with } \eta(\Sigma) \subset \mathcal{U}_f\}.$$

We've not yet assigned a topology to $W^{k,p}(\Sigma, M)$, but a topology emerges naturally from the nontrivial observation that our definition gives rise to a smooth

Banach manifold structure. Indeed, the charts are the maps $\exp_f \eta \mapsto \eta$ which take subsets of $W^{k,p}(\Sigma, M)$ into open subsets of Banach spaces, namely

$$W^{k,p}(\mathcal{U}_f) := \{\eta \in W^{k,p}(f^*TM) \mid \eta(\Sigma) \subset \mathcal{U}_f\}.$$

Since the exponential map is smooth, one can use a slight generalization of Lemma 2.67 to show that the resulting transition maps are smooth—this depends fundamentally on the same three properties of $W^{k,p}$ that were listed in the lemma: it embeds into C^0 , it is a Banach algebra, and it behaves continuously under composition with smooth functions. In the same manner, one shows that the transition maps arising from different choices of connection on M are also smooth, thus the smooth structure of $W^{k,p}(\Sigma, M)$ doesn't depend on this choice. The complete details of these arguments (in a very general context) are carried out in [Elĭ67]. The same paper also shows that the tangent spaces to $W^{k,p}(\Sigma, M)$ are canonically isomorphic to exactly what one would expect:

$$T_u W^{k,p}(\Sigma, M) = W^{k,p}(u^*TM).$$

Note that in general, $u^*TM \rightarrow \Sigma$ is only a bundle of class $W^{k,p}$, but the resulting Banach space of sections is well defined due to Exercise 3.3 above.

One can take these ideas further and speak of vector bundles whose fibers are Banach spaces: a **Banach space bundle** of class C^k is defined by a system of local trivializations whose transition maps are of class C^k from open subsets of the base to the Banach space of bounded endomorphisms $\mathcal{L}(X)$ on some Banach space X . Note that if $g : \mathcal{U} \rightarrow \mathcal{L}(X)$ is a transition map and $z \in \mathcal{U}$, $x \in X$, it is not enough to require continuity or smoothness of the map $(z, x) \mapsto g(z)x$; that is a significantly weaker condition in infinite dimensions. We refer to [Lan99] for more on the general properties of Banach space bundles.

For our purposes, it will be important to consider the Banach manifold

$$\mathcal{B}^{k,p} := W^{k,p}(\Sigma, M)$$

with a Banach space bundle $\mathcal{E}^{k-1,p} \rightarrow \mathcal{B}^{k,p}$ whose fiber at $u \in \mathcal{B}^{k,p}$ is

$$\mathcal{E}_u^{k-1,p} := W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, u^*TM)).$$

You should take a moment to convince yourself that for any $u \in \mathcal{B}^{k,p}$, it makes sense to speak of sections of class $W^{k-1,p}$ on the bundle $\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, u^*TM) \rightarrow \Sigma$. As it turns out, the general framework of [Elĭ67] implies that $\mathcal{E}^{k-1,p} \rightarrow \mathcal{B}^{k,p}$ admits the structure of a smooth Banach space bundle such that

$$\bar{\partial}_J : \mathcal{B}^{k,p} \rightarrow \mathcal{E}^{k-1,p} : u \mapsto Tu + J \circ Tu \circ j$$

is a smooth section. Note that in the last observation, we are using the assumption that J is smooth, as the question can be reduced to yet another application of Lemma 2.67: the section $\bar{\partial}_J$ contains the map $W^{k,p} \rightarrow W^{k,p} : u \mapsto J \circ u$, which has only as many derivatives as J (minus some constant). For this reason, we will assume whenever possible from now on that J is smooth.

The zero set of $\bar{\partial}_J$ is the space of solutions (3.1), and as we already observed, the topology of this solution space will have no dependence on k or p . To show that $\bar{\partial}_J^{-1}(0)$ has a nice structure, we want to apply the infinite dimensional bundle

version of the implicit function theorem, which will apply near any point $u \in \bar{\partial}_J^{-1}(0)$ at which the linearization

$$\mathbf{D}_u := D\bar{\partial}_J(u) : T_u\mathcal{B}^{k,p} \rightarrow \mathcal{E}_u^{k-1,p}$$

is surjective and has a bounded right inverse. Here \mathbf{D}_u is the operator we derived in §2.4; at the time we were assuming everything was smooth, but the result clearly extends to a bounded linear operator

$$\begin{aligned} \mathbf{D}_u : W^{k,p}(u^*TM) &\rightarrow W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, u^*TM)) \\ \eta &\mapsto \nabla\eta + J(u) \circ \nabla\eta \circ j + (\nabla_\eta J)Tu \circ j, \end{aligned}$$

where ∇ is an arbitrary symmetric connection on M , and this operator must be identical to $D\bar{\partial}_J(u)$ since C^∞ is dense in all the spaces under consideration.

The condition that \mathbf{D}_u have a bounded right inverse will turn out to be trivially satisfied whenever \mathbf{D}_u is surjective, because $\ker \mathbf{D}_u$ is finite dimensional. This is an important new feature of the global setting that did not exist locally, and we will spend the rest of this chapter proving it and computing the dimension. The main result can be summarized as follows.

Theorem 3.5. *For any $u \in \bar{\partial}_J^{-1}(0)$, \mathbf{D}_u is a Fredholm operator with index*

$$\text{ind}(\mathbf{D}_u) = n\chi(\Sigma) + 2\langle c_1(TM), [u] \rangle,$$

where $[u] := u_*[\Sigma] \in H_2(M)$ and $c_1(TM) \in H^2(M)$ is the first Chern class of the complex vector bundle (TM, J) .

Recall that a bounded linear operator $D : X \rightarrow Y$ between Banach spaces is called **Fredholm** if both $\ker D$ and $Y/\text{im } D$ are finite dimensional; the latter space is called the **cokernel** of D , often written as $\text{coker } D$. The **Fredholm index** of D is then defined to be

$$\text{ind}(D) = \dim \ker(D) - \dim \text{coker}(D).$$

Fredholm operators have many nice things in common with linear maps on finite dimensional spaces. Proofs of the following standard facts may be found in e.g. [Tay96, Appendix A] and [AA02, §4.4].

Proposition 3.6. *Assume X and Y are Banach spaces, and let $\text{Fred}(X, Y) \subset \mathcal{L}(X, Y)$ denote the space of Fredholm operators from X to Y .*

- (1) $\text{Fred}(X, Y)$ is an open subset of $\mathcal{L}(X, Y)$.
- (2) The map $\text{ind} : \text{Fred}(X, Y) \rightarrow \mathbb{Z}$ is continuous.
- (3) If $D \in \text{Fred}(X, Y)$ and $K \in \mathcal{L}(X, Y)$ is a compact operator, then $D + K \in \text{Fred}(X, Y)$.
- (4) If $D \in \text{Fred}(X, Y)$ then $\text{im } D$ is a closed subspace of Y , and there exists a closed linear subspace $V \subset X$ and finite dimensional subspace $W \subset Y$ such that

$$X = \ker(D) \oplus V, \quad Y = \text{im}(D) \oplus W,$$

and $D|_V : V \rightarrow \text{im}(D)$ is a Banach space isomorphism.

Note that the continuity of the map $\text{ind} : \text{Fred}(X, Y) \rightarrow \mathbb{Z}$ means it is locally constant, thus for any continuous family of Fredholm operators $\{D_t\}_{t \in [0,1]}$, $\text{ind}(D_t)$ is constant. This fact is extremely useful for index computations, and is true despite the fact that the dimensions of $\ker D_t$ and $Y/\text{im } D_t$ may each change quite drastically. As a simple application, this implies that for any compact operator K , $\text{ind}(D + K) = \text{ind}(D)$, as these two are connected by the continuous family $D + tK$.

Exercise 3.7. The definition of a Fredholm operator $D : X \rightarrow Y$ often includes the assumption that $\text{im } D$ is closed, but this is redundant. Convince yourself that for any $D \in \mathcal{L}(X, Y)$, if $Y/\text{im } D$ is finite dimensional then $\text{im } D$ is closed. If you get stuck, see [AA02, Corollary 2.17].

Theorem 3.5 is of course most interesting in the case where \mathbf{D}_u is surjective, as then the implicit function theorem yields:

Corollary 3.8. *If $u \in \bar{\partial}_J^{-1}(0)$ and \mathbf{D}_u is surjective, then a neighborhood of u in $\bar{\partial}_J^{-1}(0)$ admits the structure of a smooth finite dimensional manifold, with*

$$\dim \bar{\partial}_J^{-1}(0) = n\chi(\Sigma) + 2\langle c_1(TM), [u] \rangle.$$

3.2. Formal adjoints

The Fredholm theory for the operator \mathbf{D}_u fits naturally into the more general context of Cauchy-Riemann type operators on vector bundles. For the next three sections, we will consider an arbitrary smooth complex vector bundle $(E, J) \rightarrow (\Sigma, j)$ of (complex) rank n , where (Σ, j) is a closed connected Riemann surface unless otherwise noted. We will often abbreviate the first Chern number of (E, J) by writing

$$c_1(E) := \langle c_1(E, J), [\Sigma] \rangle \in \mathbb{Z}.$$

Let $D : \Gamma(E) \rightarrow \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$ denote a (real or complex) linear Cauchy-Riemann type operator. In order to understand the properties of this operator, it will be extremely useful to observe that it has a **formal adjoint**,

$$D^* : \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E)) \rightarrow \Gamma(E),$$

which will turn out to have all the same nice properties of a Cauchy-Riemann type operator. We'll use this in the next section to understand the cokernel of D , which turns out to be naturally isomorphic to the kernel of D^* .

Choose a Hermitian bundle metric $\langle \cdot, \cdot \rangle$ on E , and let (\cdot, \cdot) denote its real part, which is a real bundle metric that is invariant under the action of J . Choose also a Riemannian metric g on Σ that is compatible with the conformal structure defined by j ; this defines a volume form μ_g on Σ , and conversely (since $\dim_{\mathbb{R}} \Sigma = 2$), such a volume form uniquely determines the compatible metric g via the relation

$$\mu_g(X, Y) = g(jX, Y).$$

These choices naturally induce a bundle metric $(\cdot, \cdot)_g$ on $\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E)$, and both $\Gamma(E)$ and $\Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$ now inherit natural L^2 -inner products, defined by

$$\langle \xi, \eta \rangle_{L^2} = \int_{\Sigma} (\xi, \eta) \mu_g, \quad \langle \alpha, \beta \rangle_{L^2} = \int_{\Sigma} (\alpha, \beta)_g \mu_g$$

for $\xi, \eta \in \Gamma(E)$ and $\alpha, \beta \in \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$. We say that an operator $D^* : \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E)) \rightarrow \Gamma(E)$ is the **formal adjoint** of D if it satisfies

$$(3.3) \quad \langle \alpha, D\eta \rangle_{L^2} = \langle D^*\alpha, \eta \rangle_{L^2}.$$

for all smooth sections $\eta \in \Gamma(E)$ and $\alpha \in \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$. The existence of such operators is a quite general phenomenon that is easy to see locally using integration by parts: roughly speaking, if D has the form $D = \bar{\partial} + A$ in some local trivialization, then we expect D^* in the same local picture to take the form $-\partial + A^T$. One sees also from this local expression that D^* is almost a Cauchy-Riemann type operator; to be precise, it is *conjugate* to a Cauchy-Riemann type operator. The extra minus sign can be removed by an appropriate bundle isomorphism, and one can always transform $\partial = \partial_s - i\partial_t$ into $\bar{\partial} = \partial_s + i\partial_t$ by reversing the complex structure on the bundle. Globally, the result will be the following.

Proposition 3.9. *For any choice of Hermitian bundle metric on $(E, J) \rightarrow (\Sigma, j)$ and Riemannian metric g on Σ compatible with j , every linear Cauchy-Riemann type operator $D : \Gamma(E) \rightarrow \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$ admits a formal adjoint*

$$D^* : \Gamma(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E)) \rightarrow \Gamma(E)$$

which is conjugate to a linear Cauchy-Riemann type operator in the following sense. Defining a complex vector bundle (\hat{E}, \hat{J}) over Σ by

$$(\hat{E}, \hat{J}) := (\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E), -J),$$

there exist smooth real linear vector bundle isomorphisms

$$\Phi : \hat{E} \rightarrow \overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E), \quad \Psi : E \rightarrow \overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, \hat{E})$$

such that $\Psi \circ D^ \circ \Phi$ is a linear Cauchy-Riemann type operator on \hat{E} .*

We will prove this by deriving a global expression for D^* . One can construct it by a generalization of the same procedure by which one constructs the formal adjoint of d on the algebra of differential forms, so let us recall this first. If M is any smooth oriented manifold of real dimension m with a Riemannian metric g , let μ_g denote the induced volume form, and use g also to denote the natural extension of g to a bundle metric on each of the skew-symmetric tensor bundles $\Lambda^k T^*M$ for $k = 0, \dots, m$. We will denote $\Omega^k(M) := \Gamma(\Lambda^k T^*M)$, i.e. this is simply the vector space of smooth differential k -forms on M . Now for each $k = 0, \dots, m$, there is a unique bundle isomorphism,

$$* : \Lambda^k T^*M \rightarrow \Lambda^{m-k} T^*M$$

the **Hodge star operator**, which has the property that for all $\alpha, \beta \in \Omega^k(M)$,

$$(3.4) \quad g(\alpha, \beta) \mu_g = \alpha \wedge * \beta.$$

One can easily show that $*$ is a bundle isometry and satisfies $*^2 = (-1)^{k(m-k)}$. With this, one can associate to the exterior derivative $d : \Omega^k(M) \rightarrow \Omega^{k+1}(M)$ a formal adjoint

$$d^* : \Omega^k(M) \rightarrow \Omega^{k-1}(M),$$

$$d^* = (-1)^{m(k+1)+1} * d *,$$

which satisfies

$$\int_M g(\alpha, d\beta) \mu_g = \int_M g(d^*\alpha, \beta) \mu_g$$

for any $\alpha \in \Omega^k(M)$ and $\beta \in \Omega^{k-1}(M)$ with compact support. The proof of this relation is an easy exercise in Stokes' theorem, using (3.4).

We can extend this discussion to bundle-valued differential forms: given a real vector bundle $E \rightarrow M$, let $\Omega^k(M, E) := \Gamma(\Lambda^k T^*M \otimes E)$, which is naturally identified with the space of smooth k -multilinear antisymmetric bundle maps $TM \oplus \dots \oplus TM \rightarrow E$. Choosing a bundle metric (\cdot, \cdot) on E , the combination of g and (\cdot, \cdot) induces a natural tensor product metric on $\Lambda^k T^*M \otimes E$, which we'll denote by $(\cdot, \cdot)_g$. There is also an isomorphism of E to its dual bundle $E^* \rightarrow M$, defined by

$$E \rightarrow E^* : v \mapsto \bar{v} := (v, \cdot),$$

which extends naturally to an isomorphism

$$\Lambda^k T^*M \otimes E \rightarrow \Lambda^k T^*M \otimes E^* : \alpha \mapsto \bar{\alpha}.$$

There is no natural product structure on $\Lambda^* T^*M \otimes E$, but the wedge product does define a natural pairing

$$(\Lambda^* T^*M \otimes E^*) \oplus (\Lambda^* T^*M \otimes E) \rightarrow \Lambda^* T^*M : (\alpha \otimes \lambda, \beta \otimes v) \mapsto \lambda(v) \cdot \alpha \wedge \beta,$$

as well as a fiberwise module structure,

$$\Lambda^k T^*M \oplus (\Lambda^\ell T^*M \otimes E) \rightarrow \Lambda^{k+\ell} T^*M \otimes E : (\alpha, \beta) \mapsto \alpha \wedge \beta,$$

so that in particular $\Omega^*(M, E)$ becomes an $\Omega^*(M)$ -module.

Now if $\nabla : \Gamma(E) \rightarrow \Gamma(\text{Hom}(TM, E)) = \Omega^1(M, E)$ is a connection on $E \rightarrow M$, this has a natural extension to a **covariant exterior derivative**, which is a degree 1 linear map $d_\nabla : \Omega^*(M, E) \rightarrow \Omega^*(M, E)$ satisfying the graded Leibnitz rule

$$d_\nabla(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^k \alpha \wedge d_\nabla \beta$$

for all $\alpha \in \Omega^k(M)$ and $\beta \in \Omega^\ell(M, E)$. This also has a formal adjoint $d_\nabla^* : \Omega^*(M, E) \rightarrow \Omega^*(M, E)$, which is a linear map of degree -1 . We can write it down using a slight generalization of the Hodge star operator:

$$* : \Lambda^k T^*M \otimes E \rightarrow \Lambda^{m-k} T^*M \otimes E : \alpha \otimes v \mapsto * \alpha \otimes v,$$

in other words for any $p \in M$, $\alpha \in \Lambda^k T_p^*M$ and $v \in E_p$, the product αv defines a skew-symmetric k -form on $T_p M$ with values in E_p , and we define $*(\alpha v)$ to be $(*\alpha)v$. This map has the property that for all $\alpha, \beta \in \Omega^k(M, E)$,

$$(\alpha, \beta)_g \mu_g = \bar{\alpha} \wedge * \beta,$$

and it is then straightforward to verify that

$$(3.5) \quad \begin{aligned} d_\nabla^* : \Omega^k(M, E) &\rightarrow \Omega^{k-1}(M, E), \\ d_\nabla^* &= (-1)^{m(k+1)+1} * d_\nabla * \end{aligned}$$

has the desired property, namely that

$$(3.6) \quad \int_M (\alpha, d_\nabla \beta)_g \mu_g = \int_M (d_\nabla^* \alpha, \beta)_g \mu_g$$

for all $\alpha \in \Omega^k(M, E)$ and $\beta \in \Omega^{k-1}(M, E)$ with compact support.

Let us now extend some of these constructions to a complex vector bundle (E, J) of rank n over a complex manifold (Σ, j) of (complex) dimension m . Here it becomes natural to split the space of bundle-valued 1-forms $\Omega^1(\Sigma, E)$ into the subspaces of complex linear and antilinear forms, often called $(1, 0)$ -forms and $(0, 1)$ -forms respectively,

$$\Omega^1(\Sigma, E) = \Omega^{1,0}(\Sigma, E) \oplus \Omega^{0,1}(\Sigma, E),$$

where by definition $\Omega^{1,0}(\Sigma, E) = \Gamma(\text{Hom}_{\mathbb{C}}(T\Sigma, E))$ and $\Omega^{0,1}(\Sigma, E) = \Gamma(\overline{\text{Hom}_{\mathbb{C}}}(T\Sigma, E))$. Choosing holomorphic local coordinates (z^1, \dots, z^m) on some open subset of Σ , all the $(1, 0)$ -forms can be written on this subset as

$$\alpha = \sum_{j=1}^m \alpha_j dz^j$$

for some local sections α_j of E , and the $(0, 1)$ -forms likewise take the form

$$\alpha = \sum_{j=1}^m \alpha_j d\bar{z}^j.$$

The space of bundle-valued k -forms then splits into subspaces of (p, q) -forms for $p + q = k$,

$$\Omega^k(\Sigma, E) = \bigoplus_{p+q=k} \Omega^{p,q}(\Sigma, E),$$

where any $\alpha \in \Omega^{p,q}(\Sigma, E)$ can be written locally as a linear combination of terms of the form

$$dz^{j_1} \wedge \dots \wedge dz^{j_p} \wedge d\bar{z}^{k_1} \wedge \dots \wedge d\bar{z}^{k_q}$$

multiplied with local sections of E . The (p, q) -forms are sections of a vector bundle

$$\Lambda^{p,q} T^* \Sigma \otimes E,$$

which is a subbundle of $\Lambda^{p+q} T^* \Sigma \otimes E$.

As a special case, let $\Omega^{p,q}(\Sigma) := \Omega^{p,q}(\Sigma, \Sigma \times \mathbb{C})$ denote the space of complex valued (p, q) -forms. Then the image of the exterior derivative on $\Omega^{p,q}(\Sigma)$ splits naturally:

$$d : \Omega^{p,q}(\Sigma) \rightarrow \Omega^{p+1,q}(\Sigma) \oplus \Omega^{p,q+1}(\Sigma),$$

and with respect to this splitting we can define linear operators

$$\partial : \Omega^{p,q}(\Sigma) \rightarrow \Omega^{p+1,q}(\Sigma), \quad \bar{\partial} : \Omega^{p,q}(\Sigma) \rightarrow \Omega^{p,q+1}(\Sigma)$$

such that $d = \partial + \bar{\partial}$. The restriction to $\Omega^{0,0}(\Sigma) = C^\infty(\Sigma, \mathbb{C})$ gives (up to a factor of two)¹ the usual operators ∂ and $\bar{\partial}$ on smooth functions $f : \Sigma \rightarrow \mathbb{C}$, namely

$$\partial f = \frac{1}{2}(df - i df \circ j), \quad \bar{\partial} f = \frac{1}{2}(df + i df \circ j).$$

¹For this section only, we are modifying our usual definition of the operators $\bar{\partial}$ and ∂ on $C^\infty(\Sigma, \mathbb{C})$ to include the extra factor of $1/2$. The difference is harmless.

It follows now almost tautologically that ∂ and $\bar{\partial}$ satisfy graded Leibnitz rules,

$$\begin{aligned}\partial(\alpha \wedge \beta) &= \partial\alpha \wedge \beta + (-1)^{p+q}\alpha \wedge \partial\beta, \\ \bar{\partial}(\alpha \wedge \beta) &= \bar{\partial}\alpha \wedge \beta + (-1)^{p+q}\alpha \wedge \bar{\partial}\beta\end{aligned}$$

for $\alpha \in \Omega^{p,q}(\Sigma)$ and $\beta \in \Omega^{r,s}(\Sigma)$.

Choosing a Hermitian metric on the bundle $(E, J) \rightarrow (\Sigma, j)$, we can similarly split the derivation $d_{\nabla} : \Omega^k(\Sigma, E) \rightarrow \Omega^{k+1}(\Sigma, E)$ defined by any Hermitian connection, giving rise to complex linear operators

$$\begin{aligned}\partial_{\nabla} : \Omega^{p,q}(\Sigma, E) &\rightarrow \Omega^{p+1,q}(\Sigma, E), \\ \bar{\partial}_{\nabla} : \Omega^{p,q}(\Sigma, E) &\rightarrow \Omega^{p,q+1}(\Sigma, E)\end{aligned}$$

which satisfy similar Leibnitz rules,

$$(3.7) \quad \begin{aligned}\partial_{\nabla}(\alpha \wedge \beta) &= \partial_{\nabla}\alpha \wedge \beta + (-1)^{p+q}\alpha \wedge \partial_{\nabla}\beta, \\ \bar{\partial}_{\nabla}(\alpha \wedge \beta) &= \bar{\partial}_{\nabla}\alpha \wedge \beta + (-1)^{p+q}\alpha \wedge \bar{\partial}_{\nabla}\beta\end{aligned}$$

for $\alpha \in \Omega^{p,q}(\Sigma)$ and $\beta \in \Omega^{r,s}(\Sigma, E)$. In particular, this shows that $\bar{\partial}_{\nabla} : \Omega^{p,q}(\Sigma, E) \rightarrow \Omega^{p,q+1}(\Sigma, E)$ can be regarded as a complex linear Cauchy-Riemann type operator on the bundle $\Lambda^{p,q}T^*\Sigma \otimes E$, where we identify $\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, \Lambda^{p,q}T^*\Sigma \otimes E)$ naturally with $\Lambda^{p,q+1}T^*\Sigma \otimes E$. Restricting to $\Omega^{0,0}(\Sigma, E) = \Gamma(E)$, $\bar{\partial}_{\nabla} : \Gamma(E) \rightarrow \Omega^{0,1}(\Sigma, E)$ has the form

$$\bar{\partial}_{\nabla} = \frac{1}{2}(\nabla + J \circ \nabla \circ j).$$

We are now almost ready to write down the formal adjoint of this operator. For simplicity, we restrict to the case where Σ has complex dimension one, since this is all we need. Observe that the Hodge star then defines a bundle isomorphism of $\Lambda^1T^*\Sigma$ to itself, whose natural extension to $\Lambda^1T^*\Sigma \otimes E$ is complex linear.

Exercise 3.10.

- (1) Show that for any choice of local holomorphic coordinates $z = s + it$ on Σ , $*ds = dt$ and $*dt = -ds$.
- (2) Show that for any $\alpha \in T^*\Sigma$, $*\alpha = -\alpha \circ j$.
- (3) Show that for any $\alpha \in \Lambda^{1,0}T^*\Sigma \otimes E$, $*\alpha = -J\alpha$ and for any $\alpha \in \Lambda^{0,1}T^*\Sigma \otimes E$, $*\alpha = J\alpha$. In particular, $*$ respects the splitting $\Omega^1(\Sigma, E) = \Omega^{1,0}(\Sigma, E) \oplus \Omega^{0,1}(\Sigma, E)$.

We claim now that the formal adjoint of $\bar{\partial}_{\nabla}$ is defined by a formula analogous to the operator d_{∇}^* of (3.5), namely

$$(3.8) \quad \bar{\partial}_{\nabla}^* := - * \partial_{\nabla} * : \Omega^{0,1}(\Sigma, E) \rightarrow \Omega^0(\Sigma, E).$$

In fact, this is simply the restriction of d_{∇}^* to $\Omega^{0,1}(\Sigma, E)$, as we observe that $\bar{\partial}_{\nabla}$ maps $\Omega^{0,1}(\Sigma, E)$ to $\Omega^{0,2}(\Sigma, E)$, which is trivial since Σ has only one complex dimension. Thus the claim follows easily from (3.6) and the following exercise.

Exercise 3.11. Show that $\Lambda^{1,0}T^*\Sigma \otimes E$ and $\Lambda^{0,1}T^*\Sigma \otimes E$ are orthogonal subbundles with respect to the metric $(\ , \)_g$ on $\Lambda^1T^*\Sigma \otimes E$.

It is now easy to write down the formal adjoint of a more general Cauchy-Riemann type operator.

PROOF OF PROP. 3.9. Choosing any Hermitian connection ∇ on E , Exercise 2.17 allows us to write

$$D = \bar{\partial}_\nabla + A,$$

where $A : E \rightarrow \overline{\text{Hom}}_\mathbb{C}(T\Sigma, E)$ is a smooth real linear bundle map. (Note that Exercise 2.17 dealt only with the complex linear case, but the generalization to the real case is obvious.) Extending a well known fact from linear algebra to the context of bundles, there is a unique smooth real linear bundle map $A^T : \overline{\text{Hom}}_\mathbb{C}(T\Sigma, E) \rightarrow E$ such that

$$(\alpha, A\eta)_g = (A^T\alpha, \eta)$$

for all $z \in \Sigma$, $\eta \in E_z$ and $\alpha \in \Lambda^{0,1}T_z^*\Sigma \otimes E_z$. Then the desired operator D^* is given by

$$D^* = \bar{\partial}_\nabla^* + A^T.$$

From (3.8), we see that D^* is conjugate to an operator of the form

$$D_1 = \partial_\nabla + A_1 : \Omega^{0,1}(\Sigma, E) \rightarrow \Omega^{1,1}(\Sigma, E),$$

where $A_1 : \Lambda^{0,1}T^*\Sigma \otimes E \rightarrow \Lambda^{1,1}T^*\Sigma \otimes E$ is some smooth bundle map, i.e. a “zeroth order term.” By (3.7), this satisfies the Leibnitz rule,

$$(3.9) \quad D_1(f\alpha) = (\partial f)\alpha + fD_1\alpha$$

for all smooth functions $f : \Sigma \rightarrow \mathbb{C}$. We can turn this into the Leibnitz rule for an actual Cauchy-Riemann type operator on the bundle,

$$(\widehat{E}, \widehat{J}) = (\overline{\text{Hom}}_\mathbb{C}(T\Sigma, E), -J).$$

Indeed, the identity $\widehat{E} \rightarrow \overline{\text{Hom}}_\mathbb{C}(T\Sigma, E)$ is then a complex antilinear bundle isomorphism, and there are canonical isomorphisms

$$\Lambda^{1,1}T^*\Sigma \otimes E = \text{Hom}_\mathbb{C}(T\Sigma, \Lambda^{0,1}T^*\Sigma \otimes E) = \overline{\text{Hom}}_\mathbb{C}(T\Sigma, \widehat{E}),$$

so that D_1 is now conjugate to an operator

$$D_2 : \Gamma(\widehat{E}) \rightarrow \Gamma(\overline{\text{Hom}}_\mathbb{C}(T\Sigma, \widehat{E}))$$

which satisfies $D_2(f\beta) = (\bar{\partial}f)\beta + fD_2\beta$ due to (3.9). \square

Exercise 3.12. Show that the bundle $(\widehat{E}, \widehat{J})$, as defined in Prop. 3.9 satisfies

$$c_1(\widehat{E}) = -c_1(\Lambda^{0,1}T^*\Sigma \otimes E) = -c_1(E) - n\chi(\Sigma).$$

Remark 3.13. It’s worth noting that if (Σ, j) is a general complex manifold with a Hermitian vector bundle $(E, J) \rightarrow (\Sigma, j)$ and Hermitian connection ∇ , the resulting complex linear Cauchy-Riemann type operator

$$\bar{\partial}_\nabla : \Gamma(E) \rightarrow \Gamma(\overline{\text{Hom}}_\mathbb{C}C(T\Sigma, E))$$

does *not* necessarily define a holomorphic structure if $\dim_\mathbb{C} \Sigma \geq 2$. It turns out that the required local existence result for holomorphic sections is true if and only if the map

$$\bar{\partial}_\nabla \circ \bar{\partial}_\nabla : \Gamma(E) \rightarrow \Omega^{0,2}(\Sigma, E)$$

is zero. It’s easy to see that this condition is necessary, because if there is a holomorphic structure, then $\bar{\partial}_\nabla$ looks like the standard $\bar{\partial}$ -operator in a local holomorphic

trivialization and $\bar{\partial} \circ \bar{\partial} = 0$ on $\Omega^*(\Sigma, E)$. The converse is, in some sense, a complex version of the Frobenius integrability theorem: indeed, the corresponding statement in real differential geometry is that vector bundles with connections locally admit flat sections if and only if $d_{\nabla} \circ d_{\nabla} = 0$, which means the curvature vanishes. A proof of the complex version may be found in [DK90, § 2.2.2], and the first step in this proof is the local existence result for the case $\dim_{\mathbb{C}} \Sigma = 1$ (our Theorem 2.49). Observe that the integrability condition is trivially satisfied when $\dim_{\mathbb{C}} \Sigma = 1$, since then $\Omega^{0,2}(\Sigma, E)$ is a trivial space.

3.3. The Fredholm property

For the remainder of this chapter, (Σ, j) will be a closed Riemann surface and $(E, J) \rightarrow (\Sigma, j)$ will be a complex vector bundle of rank n with a real linear Cauchy-Riemann operator D . We shall now prove the Fredholm property for the obvious extension of D to a bounded linear map

$$(3.10) \quad D : W^{k,p}(E) \rightarrow W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E)),$$

with $k \in \mathbb{N}$ and $p \in (1, \infty)$.

Theorem 3.14. *The operator D of (3.10) is Fredholm, and neither $\ker D$ nor $\text{ind}(D)$ depends on the choice of k and p .*

This result depends essentially on three ingredients: first, the Calderón-Zygmund inequality gives an estimate for $\|\eta\|_{W^{k,p}}$ in terms of $\|D\eta\|_{W^{k-1,p}}$, from which we will be able to show quite easily that $\ker D$ is finite dimensional. The second ingredient is the formal adjoint D^* that was derived in the previous section: since D^* is also conjugate to a Cauchy-Riemann type operator, the previous step implies that its kernel is also finite dimensional. The final ingredient is elliptic regularity, which we can use to identify the cokernel of D with the kernel of D^* . The regularity theory also implies that both of these kernels consist only of smooth sections, and are thus completely independent of k and p .

As sketched above, the first step in proving Theorem 3.14 is an a priori estimate that follows from the linear regularity theory of §2.5. In particular, the Calderón-Zygmund inequality (Theorem 2.35 and Exercise 2.36) gives

$$(3.11) \quad \|\eta\|_{W^{k,p}} \leq c\|\eta\|_{W^{k-1,p}} \text{ for all } \eta \in W_0^{k,p}(B, \mathbb{C}^n).$$

If $\eta \in W^{k,p}(E)$ and $\{\alpha_j\}$ is a partition of unity subordinate to some finite open cover of Σ , then we can apply (3.11) in local trivializations and charts to the sections $\alpha_i \eta$. In this local picture, D becomes an operator of the form $\bar{\partial} + A$, where A is a smooth family of matrices, thus locally defining a bounded linear operator from $W^{k,p}$ to itself. The result is the following estimate, of which a more detailed proof may be found in [MS04, Lemma C.2.1].

Lemma 3.15. *For each $k \in \mathbb{N}$ and $p \in (1, \infty)$, there exists a constant $c > 0$ such that for every $\eta \in W^{k,p}(E)$,*

$$\|\eta\|_{W^{k,p}(E)} \leq c\|D\eta\|_{W^{k-1,p}(E)} + c\|\eta\|_{W^{k-1,p}(E)}.$$

Observe that the inclusion $W^{k,p}(E) \hookrightarrow W^{k-1,p}(E)$ is compact. This will allow us to make use of the following general result.

Proposition 3.16. *Suppose X , Y and Z are Banach spaces, $A \in \mathcal{L}(X, Y)$, $K \in \mathcal{L}(X, Z)$ is compact, and there is a constant $c > 0$ such that for all $x \in X$,*

$$(3.12) \quad \|x\|_X \leq c\|Ax\|_Y + c\|Kx\|_Z.$$

Then $\ker A$ is finite dimensional and $\operatorname{im} A$ is closed.

PROOF. A vector space is finite dimensional if and only if the unit ball in that space is a compact set, so we begin by proving the latter holds for $\ker A$. Suppose $x_k \in \ker A$ is a bounded sequence. Then since K is a compact operator, Kx_k has a convergent subsequence in Z , which is therefore Cauchy. But (3.12) then implies that the corresponding subsequence of x_k in X is also Cauchy, and thus converges.

Since we now know $\ker A$ is finite dimensional, we also know there is a closed complement $V \subset X$ with $\ker A \oplus V = X$. Then the restriction $A|_V$ has the same image as A , thus if $y \in \overline{\operatorname{im} A}$, there is a sequence $x_k \in V$ such that $Ax_k \rightarrow y$. We claim that x_k is bounded. If not, then $A(x_k/\|x_k\|_X) \rightarrow 0$ and $K(x_k/\|x_k\|_X)$ has a convergent subsequence, so (3.12) implies that a subsequence of $x_k/\|x_k\|_X$ also converges to some $x_\infty \in V$ with $\|x_\infty\| = 1$ and $Ax_\infty = 0$, a contradiction. But now since x_k is bounded, Kx_k also has a convergent subsequence and Ax_k converges by assumption, thus (3.12) yields also a convergent subsequence of x_k , whose limit x satisfies $Ax = y$. This completes the proof that $\operatorname{im} A$ is closed. \square

The above implies that every Cauchy-Riemann type operator has finite dimensional kernel and closed image; operators with these two properties are called **semi-Fredholm**. Note that by elliptic regularity, $\ker D$ only contains smooth sections, and is thus the same space for every k and p .

By Prop. 3.9, the same results obviously apply to the formal adjoint, after extending it to a bounded linear operator

$$D^* : W^{k,p}(\overline{\operatorname{Hom}}_{\mathbb{C}}(T\Sigma, E)) \rightarrow W^{k-1,p}(E).$$

Proposition 3.17. *Using the natural inclusion $W^{k,p} \hookrightarrow W^{k-1,p}$ to inject $\ker D$ and $\ker D^*$ into $W^{k-1,p}$, there are direct sum splittings*

$$\begin{aligned} W^{k-1,p}(\overline{\operatorname{Hom}}_{\mathbb{C}}(T\Sigma, E)) &= \operatorname{im} D \oplus \ker D^* \\ W^{k-1,p}(E) &= \operatorname{im} D^* \oplus \ker D. \end{aligned}$$

Thus the projections along $\operatorname{im} D$ and $\operatorname{im} D^$ yield natural isomorphisms $\operatorname{coker} D = \ker D^*$ and $\operatorname{coker} D^* = \ker D$.*

PROOF. We will prove only the first of the two splittings, as the second is entirely analogous. We claim first that $\operatorname{im} D \cap \ker D^* = \{0\}$. Indeed, if $\alpha \in W^{k-1,p}(\overline{\operatorname{Hom}}_{\mathbb{C}}(T\Sigma, E))$ with $D^*\alpha = 0$, then since D^* is conjugate to a Cauchy-Riemann type operator via smooth bundle isomorphisms, elliptic regularity implies that α is smooth. Then if $\alpha = D\eta$ for some $\eta \in W^{k,p}(E)$, η must also be smooth, and we find

$$0 = \langle D^*\alpha, \eta \rangle_{L^2} = \langle \alpha, D\eta \rangle_{L^2} = \|\alpha\|_{L^2}^2.$$

To show that $\text{im } D + \ker D^* = W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$, it will be convenient to address the case $k = 1$ first. Note that $\text{im } D + \ker D^*$ is a closed subspace since $\text{im } D$ is closed and $\ker D^*$ is finite dimensional. Then if it is not all of L^p , there exists a nonzero $\alpha \in L^q(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$, where $\frac{1}{p} + \frac{1}{q} = 1$, such that

$$\begin{aligned}\langle \alpha, D\eta \rangle_{L^2} &= 0 \text{ for all } \eta \in W^{1,p}(E), \\ \langle \alpha, \beta \rangle_{L^2} &= 0 \text{ for all } \beta \in \ker D^*.\end{aligned}$$

The first relation is valid in particular for all smooth η , and this means that α is a weak solution of the equation $D^*\alpha = 0$, so by regularity of weak solutions (see Corollary 2.46), α is smooth and belongs to $\ker D^*$. Then we can plug $\beta = \alpha$ into the second relation and conclude $\alpha = 0$.

Now we show that $\text{im } D + \ker D^* = W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$ when $k \geq 2$. Given $\alpha \in W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$, α is also of class L^p and thus the previous step gives $\eta \in W^{1,p}(E)$ and $\beta \in \ker D^*$ such that

$$D\eta + \beta = \alpha.$$

Then β is smooth, and $D\eta = \alpha - \beta$ is of class $W^{k-1,p}$, so regularity (Corollary 2.46 again) implies that $\eta \in W^{k,p}(E)$, and we are done. \square

We are now finished with the proof of Theorem 3.14, as we have shown that both $\ker D$ and $\ker D^* \cong \text{coker } D$ are finite dimensional spaces consisting only of smooth sections, which are thus contained in $W^{k,p}$ for all k and p .

Exercise 3.18. This exercise is meant to convince you that “boundary conditions are important.” Recall that the Calderón-Zygmund inequality $\|u\|_{W^{1,p}} \leq c\|\bar{\partial}u\|_{L^p}$ is valid for smooth \mathbb{C}^n -valued functions u with compact support in the open unit ball $B \subset \mathbb{C}$. Show that this inequality cannot be extended to functions without compact support; in fact there is not even any estimate of the form

$$\|u\|_{W^{1,p}} \leq c\|\bar{\partial}u\|_{L^p} + c\|u\|_{L^p}$$

for general functions $u \in C^\infty(B) \cap W^{1,p}(B)$. Why not? For contrast, see Exercise 3.23 below.

3.4. The Riemann-Roch formula and transversality criteria

It is easy to see that the index of a Cauchy-Riemann type operator $D : W^{k,p}(E) \rightarrow W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$ depends only on the isomorphism class of the bundle $(E, J) \rightarrow (\Sigma, j)$. Indeed, by Exercise 2.17, the difference between any two such operators D and D' on the same bundle defines a smooth real linear bundle map $A : E \rightarrow \overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E)$ such that

$$D'\eta - D\eta = A\eta.$$

We often refer to this bundle map as a “zeroth order term.” It defines a bounded linear map from $W^{k,p}(E)$ to $W^{k,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$, which is then composed with the compact inclusion into $W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, E))$ and is therefore a compact operator.

We conclude that all Cauchy-Riemann type operators on the same bundle are compact perturbations of each other,² and thus have the same Fredholm index. Since complex vector bundles over a closed surface are classified up to isomorphism by the first Chern number, the index will therefore depend only on the topological type of Σ and on $c_1(E)$. To compute it, we can use the fact that every complex bundle admits a *complex linear* Cauchy-Riemann operator (cf. Exercise 2.18), and restrict our attention to the complex linear case. Then E is a holomorphic vector bundle, and $\ker D$ is simply the vector space of holomorphic sections. We'll see below that in some important examples, it is not hard to compute this space explicitly. The key observation is that one can identify holomorphic sections on vector bundles over Σ with complex valued meromorphic functions on Σ that have prescribed poles and/or zeroes. The problem of understanding such spaces of meromorphic functions is a classical one, and its solution is the Riemann-Roch formula.

Theorem 3.19 (Riemann-Roch formula). $\operatorname{ind}(D) = n\chi(\Sigma) + 2c_1(E)$.

We should emphasize, especially for readers who are more accustomed to algebraic geometry, that this is the *real* index, i.e. the difference between $\dim \ker D$ and $\dim \operatorname{coker} D$ as real vector spaces—these dimensions may indeed be odd in general since we'll be interested in cases where D is not complex linear, but $\operatorname{ind}(D)$ will always be even, a nontrivial consequence of the fact that D is always *homotopic* to a complex linear operator. We will later see cases (on punctured Riemann surfaces or surfaces with boundary) where $\operatorname{ind}(D)$ can also be odd.

A complete proof of the Riemann-Roch formula may be found in [MS04, Appendix C] or, from a more classical perspective, any number of books on Riemann surfaces. Below we will explain a proof for the genus 0 case and give a heuristic argument to justify the rest. An important feature will be the following “transversality” criterion, which will also have many important applications in the study of J -holomorphic curves. It is a consequence of the identification $\ker D \equiv \operatorname{coker} D^*$, combined with the similarity principle (recall §2.7).

Theorem 3.20. *Suppose $n = 1$, i.e. $(E, J) \rightarrow (\Sigma, j)$ is a complex line bundle.*

- *If $c_1(E) < 0$, then D is injective.*
- *If $c_1(E) > -\chi(\Sigma)$, then D is surjective.*

PROOF. The criterion for injectivity is an easy consequence of the similarity principle, for which we don't really need to know anything about D except that it's a Cauchy-Riemann type operator. If $E \rightarrow \Sigma$ has complex rank 1 and $\ker D$ contains a nontrivial section η , then by the similarity principle, η has only isolated (and thus finitely many) zeroes, each of which counts with positive order. The count of these computes the first Chern number of E , thus $c_1(E) \geq 0$, and D must be injective if $c_1(E) < 0$.

²This statement is false when Σ is not compact: we'll see when we later discuss Cauchy-Riemann type operators on domains with cylindrical ends that the zeroth order term is no longer compact, and the index does depend on the behavior of this term at infinity.

The second part follows now from the observation that D is surjective if and only if D^* is injective, and the latter is guaranteed by the condition $c_1(\widehat{E}) < 0$, which by Prop. 3.9 and Exercise 3.12 is equivalent to $c_1(E) > -\chi(\Sigma)$. \square

Observe that we did not need to know the index formula in order to deduce the last result. In fact, this already gives enough information to deduce the index formula in the special case $\Sigma = S^2$, which will be the most important in our applications.

PROOF OF THEOREM 3.19 IN THE CASE $\Sigma = S^2$. We assume first that $n = 1$. In this situation, at least one of the criteria $c_1(E) < 0$ or $c_1(E) > -\chi(\Sigma) = -2$ from Theorem 3.20 is always satisfied, hence D is always injective or surjective; in fact if $c_1(E) = -1$ it is an isomorphism. By considering D^* instead of D if necessary, we can restrict our attention to the case where D is surjective, so $\text{ind } D = \dim \ker D$. We will now construct for each value of $c_1(E) \geq 0$ a “model” holomorphic line bundle, which is sufficiently simple so that we can identify the space of holomorphic sections explicitly.

For the case $c_1(E) = 0$, the model bundle is obvious: just take the trivial line bundle $S^2 \times \mathbb{C} \rightarrow S^2$, so the holomorphic sections are holomorphic functions $S^2 \rightarrow \mathbb{C}$, which are necessarily constant and therefore $\dim \ker D = 2$, as it should be. A more general model bundle can be defined by gluing together two local trivializations: let $E^{(1)}$ and $E^{(2)}$ denote two copies of the trivial holomorphic line bundle $\mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$, and for any $k \in \mathbb{Z}$, define

$$E_k := (E^{(1)} \sqcup E^{(2)}) / (z, v) \sim \Phi_k(z, v),$$

where $\Phi_k : E^{(1)}|_{\mathbb{C} \setminus \{0\}} \rightarrow E^{(2)}|_{\mathbb{C} \setminus \{0\}}$ is a bundle isomorphism covering the biholomorphic map $z \mapsto 1/z$ and defined by $\Phi_k(z, v) = (1/z, g_k(z)v)$, with

$$g_k(z)v := \frac{1}{z^k}v.$$

The function $g_k(z)$ is a holomorphic transition map, thus E_k has a natural holomorphic structure. Regarding a function $f : \mathbb{C} \rightarrow \mathbb{C}$ as a section of $E^{(1)}$, we have

$$\Phi_k(1/z, f(1/z)) = (z, z^k f(1/z)),$$

which means that f extends to a smooth section of E_k if and only if the function $g(z) = z^k f(1/z)$ extends smoothly to $z = 0$. It follows that $c_1(E_k) = k$, as one can choose $f(z) = 1$ for z in the unit disk and then modify $g(z) = z^k$ to a smooth function that algebraically has k zeroes at 0 (note that an actual modification is necessary only if $k < 0$). Similarly, the holomorphic sections of E_k can be identified with the entire functions $f : \mathbb{C} \rightarrow \mathbb{C}$ such that $z^k f(1/z)$ extends holomorphically to $z = 0$; if $k < 0$ this implies $f \equiv 0$, and if $k \geq 0$ it means $f(z)$ is a polynomial of degree at most k , hence $\dim \ker D = 2 + 2k$. The proof of the index formula for $\Sigma = S^2$ and $n = 1$ is now complete.

The case $n \geq 2$ can easily be derived from the above. It suffices to prove that $\text{ind}(D) = 2n + 2c_1(E)$ for some model holomorphic bundle of rank n with a given value of $c_1(E)$. Indeed, for any $k \in \mathbb{Z}$, take E to be the direct sum of n holomorphic line bundles,

$$E := E_{-1} \oplus \dots \oplus E_{-1} \oplus E_k,$$

which has $c_1(E) = k - (n - 1)$. By construction, the natural Cauchy-Riemann operator D on E splits into a direct sum of Cauchy-Riemann operators on its summands, and it is an isomorphism on each of the E_{-1} factors, thus we conclude as in the line bundle case that D is injective if $k < 0$ and surjective if $k \geq 0$. By replacing D with D^* if necessary, we can now assume without loss of generality that $k \geq 0$ and D is surjective. The space of holomorphic sections is then simply the direct sum of the corresponding spaces for its summands, which are trivial for E_{-1} and have dimension $2 + 2k$ for E_k . We therefore have

$$\operatorname{ind}(D) = \dim \ker D = 2 + 2k = 2n + 2[k - (n - 1)] = n\chi(\Sigma) + 2c_1(E).$$

□

One should not conclude from the above proof that every Cauchy-Riemann type operator on the sphere is either injective or surjective, which is true on line bundles but certainly not for bundles of higher rank—above we only used the fact that for every value of $c_1(E)$, *one can construct* a bundle that has this property. The proof is not so simple for general Riemann surfaces because it is less straightforward to identify spaces of holomorphic sections. One lesson to be drawn from the above argument, however, is that holomorphic sections on a line bundle with $c_1(E) = k$ can also be regarded as holomorphic sections on some related bundle with $c_1(E) = k + 1$, but with an extra zero at some chosen point. This suggests that an increment in the value of $c_1(E)$ should also enlarge the space of holomorphic sections by two real dimensions, because one can add two linearly independent sections that do not vanish at the chosen point. What's true for line bundles in this sense is also true for bundles of higher rank, because one can always construct model bundles that are direct sums of line bundles. We will not attempt to make this argument precise, but it should give some motivation to believe that $\operatorname{ind}(D)$ scales with $2c_1(E)$: to be exact, there exists a constant $C = C(\Sigma, n)$ such that

$$\operatorname{ind}(D) = C(\Sigma, n) + 2c_1(E).$$

If you believe this, then we can already deduce the general Riemann-Roch formula by comparing D with its formal adjoint. Indeed, D^* has index

$$\operatorname{ind}(D^*) = C(\Sigma, n) + 2c_1(\widehat{E}) = C(\Sigma, n) - 2c_1(E) - 2n\chi(\Sigma)$$

according to Exercise 3.12, and since $\operatorname{coker} D = \ker D^*$ and vice versa, $\operatorname{ind}(D^*) = -\operatorname{ind}(D)$. Thus adding these formulas together yields

$$0 = 2C(\Sigma, n) - 2n\chi(\Sigma).$$

We conclude $C(\Sigma, n) = n\chi(\Sigma)$, and the Riemann-Roch formula follows.

With the index formula understood, we can derive some alternative formulations of the transversality criteria in Theorem 3.20 which will often be useful. First, compare the formulas for $\operatorname{ind}(D)$ and $\operatorname{ind}(D^*)$:

$$\begin{aligned} \operatorname{ind}(D) &= \chi(\Sigma) + 2c_1(E), \\ \operatorname{ind}(D^*) &= \chi(\Sigma) + 2c_1(\widehat{E}), \end{aligned}$$

where \widehat{E} is the line bundle constructed in the proof of Prop. 3.9. Since $\text{ind}(D) = -\text{ind}(D^*)$, subtracting the second formula from the first yields

$$\text{ind}(D) = c_1(E) - c_1(\widehat{E}),$$

and thus $c_1(\widehat{E}) < 0$ if and only if $\text{ind}(D) > c_1(E)$, which implies by Theorem 3.20 that D^* is injective and thus D is surjective. We state this as a corollary.

Corollary 3.21. *If $n = 1$ and $\text{ind}(D) > c_1(E)$, then D is surjective.*

Exercise 3.22. Show that another equivalent formulation of Theorem 3.20 for Cauchy-Riemann operators on complex line bundles is the following:

- If $\text{ind}(D) < \chi(\Sigma)$ then D is injective.
- If $\text{ind}(D) > -\chi(\Sigma)$ then D is surjective.

This means that for line bundles, D is always surjective (or injective) as soon as its index is large (or small) enough. Observe that when $\Sigma = S^2$, one of these conditions is always satisfied, but there is always an “interval of uncertainty” in the higher genus case.

A different approach to the proof of Riemann-Roch, which is taken in [MS04], is to cut up $E \rightarrow \Sigma$ into simpler pieces on which the index can be computed explicitly, and then conclude the general result by a “linear gluing argument”. We’ll come back to this idea in a later chapter when we discuss the generalization of the Riemann-Roch formula to open surfaces with cylindrical ends. The proof in [MS04] instead considers Cauchy-Riemann operators on surfaces with boundary and totally real boundary conditions: the upshot is that the problem can be reduced in this way to the following exercise, in which one computes the index for the standard Cauchy-Riemann operator on a closed disk.

Exercise 3.23. Let $\mathbb{D} \subset \mathbb{C}$ denote the closed unit disk and E the trivial bundle $\mathbb{D} \times \mathbb{C} \rightarrow \mathbb{D}$. For a given integer $\mu \in \mathbb{Z}$, define a real rank 1 subbundle $\ell_\mu \subset E|_{\partial\mathbb{D}}$ by

$$(\ell_\mu)_{e^{i\theta}} = e^{i\pi\mu\theta} \mathbb{R} \subset \mathbb{C}.$$

We call ℓ_μ in this context a **totally real** subbundle of $\mathbb{D} \times \mathbb{C}$ at the boundary, and the integer μ is its **Maslov index**. Let $\bar{\partial} = \partial_s + i\partial_t$, and for $k, p > 2$ consider the operator

$$\bar{\partial} : W_{\ell_\mu}^{k,p}(\mathbb{D}, \mathbb{C}) \rightarrow W^{k-1,p}(\mathbb{D}, \mathbb{C}),$$

where the domain is defined by

$$W_{\ell_\mu}^{k,p}(\mathbb{D}, \mathbb{C}) = \{\eta \in W^{k,p}(\mathbb{D}, \mathbb{C}) \mid \eta(\partial\mathbb{D}) \subset \ell_\mu\}.$$

Show that as an operator between these particular spaces, $\ker \bar{\partial}$ has dimension $1 + \mu = \chi(\mathbb{D}) + \mu$ if $\mu \geq -1$, and $\bar{\partial}$ is injective if $\mu \leq -1$. (You may find it helpful to think in terms of Fourier series.) By constructing the appropriate formal adjoint of $\bar{\partial}$ in this setting (which will also satisfy a totally real boundary condition), one can also show that $\bar{\partial}$ is surjective if $\mu \geq -1$, and one can similarly compute the kernel of the formal adjoint if $\mu \leq -1$, concluding that $\bar{\partial}$ is in fact Fredholm and has index

$\text{ind}(\bar{\partial}) = \chi(\mathbb{D}) + \mu$. By considering direct sums of line bundles with totally real boundary conditions, this generalizes easily to bundles of general rank $n \in \mathbb{N}$ as

$$\text{ind}(\bar{\partial}) = n\chi(\mathbb{D}) + \mu.$$

One should think of this as another instance of the Riemann-Roch formula, in which the Maslov index now plays the role of $2c_1(E)$. The details are carried out in [MS04, Appendix C].

As a final remark, we note that the Fredholm theory of Cauchy-Riemann operators gives a new proof of a local regularity result that we made much use of in Chapter 2: the standard $\bar{\partial}$ -operator on the open unit ball $B \subset \mathbb{C}$,

$$\bar{\partial} : W^{k,p}(B, \mathbb{C}^n) \rightarrow W^{k-1,p}(B, \mathbb{C}^n)$$

has a bounded right inverse (see Theorem 2.43). This follows from our proof of Theorem 3.19 in the case $\Sigma = S^2$, because any $f \in W^{k-1,p}(B, \mathbb{C}^n)$ can be extended to a section in $W^{k-1,p}(\overline{\text{Hom}}_{\mathbb{C}}(TS^2, S^2 \times \mathbb{C}^n))$, and we can then use the fact that the standard Cauchy-Riemann operator on the trivial bundle $S^2 \times \mathbb{C}^n$ is a surjective Fredholm operator, its kernel consisting of the constant sections. Alternatively, one can use the fact established by Exercise 3.23, that the restriction of $\bar{\partial}$ to the domain $W_{\ell_0}^{k,p}(\mathbb{D}, \mathbb{C}^n)$ of functions with the totally real boundary condition $\eta(\partial\mathbb{D}) \subset \mathbb{R}^n$ is a surjective Fredholm operator with index n ; its kernel is again the space of constant functions.

CHAPTER 4

Moduli spaces

4.1. The moduli space of closed J -holomorphic curves

In the previous chapter we considered the local structure of the space of J -holomorphic maps $(\Sigma, j) \rightarrow (M, J)$ from a fixed closed Riemann surface to a fixed smooth almost complex manifold of dimension $2n$. From a geometric point of view, this is not the most natural space to study: geometrically we prefer to picture holomorphic curves as 2-dimensional submanifolds¹ whose tangent spaces are invariant under the action of J . In the symplectic context in particular, this means they give rise to symplectic submanifolds. From this perspective, the interesting object is not the parametrization u but its image $u(\Sigma)$, thus we should regard all reparametrizations of u to be equivalent. Moreover, the choice of parametrization fully determines $j = u^*J$, thus one cannot choose j in advance, but must allow it to vary over the space of all complex structures on Σ . The interesting solution space is therefore the following.

Definition 4.1. For any integers $g, m \geq 0$, and a homology class $A \in H_2(M)$, we define the **moduli space of J -holomorphic curves in M with genus g and m marked points representing A** to be

$$\mathcal{M}_{g,m}^A(J) = \{(\Sigma, j, u, (z_1, \dots, z_m))\} / \sim,$$

where (Σ, j) is any closed Riemann surface of genus g , $u : (\Sigma, j) \rightarrow (M, J)$ is a pseudoholomorphic map with $[u] := u_*[\Sigma] = A$, and (z_1, \dots, z_m) is an ordered set of distinct points in Σ , which we'll often denote by

$$\Theta = (z_1, \dots, z_m).$$

We say $(\Sigma, j, u, \Theta) \sim (\Sigma', j', u', \Theta')$ if and only if there exists a biholomorphic diffeomorphism $\varphi : (\Sigma, j) \rightarrow (\Sigma', j')$ such that $u = u' \circ \varphi$ and $\varphi(\Theta) = \Theta'$ with the ordering preserved.

The presence of the marked points $\Theta = (z_1, \dots, z_m)$ in this definition may at first seem superfluous, but we'll find them useful because there exist natural **evaluation maps**

$$\text{ev}_i : \mathcal{M}_{g,m}^A(J) \rightarrow M : [(\Sigma, j, u, \Theta)] \mapsto u(z_i),$$

which yield relations between the topology of M and the topology of the moduli space. This will be crucial in later applications to symplectic geometry.

¹This description is of course only strictly correct for holomorphic curves that are embedded, which they need not be in general—though we'll see that in many important applications, they are.

A topology on $\mathcal{M}_{g,m}^A(J)$ can be defined via the following notion of convergence: we say $[(\Sigma_k, j_k, u_k, \Theta_k)] \rightarrow [(\Sigma, j, u, \Theta)]$ if for sufficiently large k the sequence has representatives of the form $(\Sigma, j'_k, u'_k, \Theta)$ such that $j'_k \rightarrow j$ and $u'_k \rightarrow u$ in the C^∞ -topology. In particular Σ_k must be diffeomorphic to Σ and have the same number of marked points for sufficiently large k ; observe that when this is the case, one can always choose a diffeomorphism to fix the positions of the marked points.

We'll often abbreviate

$$\mathcal{M}(J) = \bigcup_{g,m,A} \mathcal{M}_{g,m}^A(J),$$

and abuse notation by writing an equivalence class of tuples $[(\Sigma, j, u, \Theta)]$ simply as (Σ, j, u, Θ) or $u \in \mathcal{M}(J)$ when there's no ambiguity. For a given $u \in \mathcal{M}(J)$, we'll also denote by

$$\mathcal{M}_u(J) \subset \mathcal{M}(J)$$

the connected component of $\mathcal{M}(J)$ that contains u . The main goal of this chapter will be to show that under suitable hypotheses, a certain subset of $\mathcal{M}_u(J)$ is a smooth finite dimensional manifold. Its “expected” or *virtual* dimension is essentially a Fredholm index with some correction terms, and depends on the topological data g , m and A . We'll use the convenient abbreviation,

$$c_1(A) = \langle c_1(TM, J), A \rangle.$$

Definition 4.2. For $u \in \mathcal{M}_{g,m}^A(J)$, define the **index of u** to be the integer

$$(4.1) \quad \text{ind}(u) = (n-3)(2-2g) + 2c_1(A) + 2m.$$

This is also called the **virtual dimension** of the moduli space near u , and denoted by $\text{vir-dim } \mathcal{M}_{g,m}^A(J)$.

It is both interesting and important to consider the special case where M is a single point: then $\mathcal{M}_{g,m}^A(J)$ reduces to the **moduli space of Riemann surfaces** with genus g and m marked points:

$$\mathcal{M}_{g,m} = \{(\Sigma, j, (z_1, \dots, z_m))\} / \sim,$$

with the equivalence and topology defined the same as above (all statements involving the map u are now vacuous). The elements $(\Sigma, j, \Theta) \in \mathcal{M}_{g,m}$ are called **pointed Riemann surfaces**, and each comes with an **automorphism group**

$$\text{Aut}(\Sigma, j, \Theta) = \{\varphi : (\Sigma, j) \rightarrow (\Sigma, j) \text{ biholomorphic} \mid \varphi|_\Theta = \text{Id}\}.$$

Similarly, a J -holomorphic curve $(\Sigma, j, u, \Theta) \in \mathcal{M}(J)$ has an automorphism group

$$\text{Aut}(u) := \text{Aut}(\Sigma, j, \Theta, u) := \{\varphi \in \text{Aut}(\Sigma, j, \Theta) \mid u = u \circ \varphi\}.$$

It turns out that in understanding the local structure of $\mathcal{M}(J)$, a special role is played by holomorphic curves with trivial automorphism groups. The following simple result was proved as Theorem 2.91 in Chapter 2, and it implies (via Exercise 4.5 below) that whenever any nontrivial holomorphic curves exist, one can also find curves with this property.

Proposition 4.3. *For any closed, connected and nonconstant J -holomorphic curve $u : (\Sigma, j) \rightarrow (M, J)$, there exists a factorization $u = v \circ \varphi$ where*

- $v : (\Sigma', j') \rightarrow (M, J)$ is a closed J -holomorphic curve that is embedded outside a finite set of critical points and self-intersections, and
- $\varphi : (\Sigma, j) \rightarrow (\Sigma', j')$ is a holomorphic map of degree $\deg(\varphi) \geq 1$.

Moreover, v is unique up to biholomorphic reparametrization.

Definition 4.4. The degree of $\varphi : \Sigma \rightarrow \Sigma'$ in Prop. 4.3 is called the **covering number** or **covering multiplicity** of u . If this is 1, then we say u is **simple**.

We denote by

$$\mathcal{M}^*(J) \subset \mathcal{M}(J)$$

the subset consisting of all curves in $\mathcal{M}(J)$ that are simple, and if $u \in \mathcal{M}^*(J)$, denote by $\mathcal{M}_u^*(J) \subset \mathcal{M}^*(J)$ the connected component containing u . By Prop. 4.3, $\mathcal{M}^*(J)$ and $\mathcal{M}_u^*(J)$ are clearly open subsets of $\mathcal{M}(J)$ and $\mathcal{M}_u(J)$ respectively.

Exercise 4.5. Show that if $u : (\Sigma, j) \rightarrow (M, J)$ is simple then $\text{Aut}(u)$ is trivial (for any choice of marked points). *Hint: it suffices to observe that a simple curve is necessarily somewhere injective, which means there exists a point $z \in \Sigma$ such that $du(z) \neq 0$ and $u^{-1}(u(z)) = \{z\}$. Show that no such point can exist if $\text{Aut}(u)$ contains a nontrivial map.*

Exercise 4.6. Show that if $u : (\Sigma, j) \rightarrow (M, J)$ has covering multiplicity $k \in \mathbb{N}$ then for any set of marked points Θ , the order of $\text{Aut}(\Sigma, j, \Theta, u)$ is at most k .

Recall that a subset Y in a complete metric space X is called a **Baire set** or said to be of **second category** if it is a countable intersection of open dense sets. The Baire category theorem implies that such subsets are also dense, and Baire sets are often used to define an infinite dimensional version of the term “almost everywhere,” i.e. they are analogous to sets whose complements have Lebesgue measure zero. It is common to say that a property is satisfied for **generic** choices of data if the set of all possible data contains a Baire subset for which the property is satisfied.

Since it is important for applications, we shall assume in the following that M carries a symplectic structure ω , and focus our attention on the space of ω -compatible almost complex structures $\mathcal{J}(M, \omega)$ that was defined in §2.2. Similarly, we continue to denote the space of ω -tame complex structures by $\mathcal{J}^\tau(M, \omega)$, and the space of all almost complex structures by $\mathcal{J}(M)$. All of these spaces carry natural C^∞ -topologies and are Fréchet manifolds if M is closed. In the following sections we will prove several results which, taken together, imply the following local structure theorem.

Theorem 4.7. *If (M, ω) is a closed symplectic manifold, then there exists a Baire set $\mathcal{J}_{\text{reg}}(M, \omega) \subset \mathcal{J}(M, \omega)$ such that for every $J \in \mathcal{J}_{\text{reg}}(M, \omega)$, the space of simple J -holomorphic curves $\mathcal{M}^*(J)$ admits the structure of a smooth finite dimensional manifold, whose dimension in $\mathcal{M}^*(J) \cap \mathcal{M}_{g,m}^A(J)$ for any $g, m \geq 0$ and $A \in H_2(M)$ is precisely the virtual dimension of $\mathcal{M}_{g,m}^A(J)$.*

Remark 4.8. The above theorem and all related results (e.g. Theorem 4.10 below) remain true if $\mathcal{J}(M, \omega)$ is replaced by $\mathcal{J}^\tau(M, \omega)$ or $\mathcal{J}(M)$; in fact, the equivalence of these last two variations is obvious since $\mathcal{J}^\tau(M, \omega)$ is an open subset of $\mathcal{J}(M)$.

Of course for $\mathcal{J}(M)$, the manifold M need not admit a symplectic structure, though one must then confront the possibility that $\mathcal{J}(M)$ might be empty. The symplectic structure will in any case play no role whatsoever in the proofs of these results except to make one detail slightly harder (see Lemma 4.48), thus it will be immediate that minor alterations of the same proofs imply the same results for tame or general almost complex structures.

One of the important consequences of Theorem 4.7 is that for generic choices of J , every connected component of the moduli space $\mathcal{M}^*(J)$ must have nonnegative virtual dimension, as a smooth manifold of negative dimension is empty by definition. Put another way, if a simple curve of negative index exists, then one can always eliminate it by a small perturbation of J . We state this as a corollary.

Corollary 4.9. *For $J \in \mathcal{J}_{\text{reg}}(M, \omega)$, every simple $u \in \mathcal{M}^*(J)$ has $\text{ind}(u) \geq 0$.*

An important related problem is to consider the space of J_τ -holomorphic curves, where $\{J_\tau\}$ is a smooth homotopy of almost complex structures. Given $J_0, J_1 \in \mathcal{J}(M, \omega)$, define

$$\mathcal{J}(M, \omega; J_0, J_1)$$

to be the space of all smooth 1-parameter families $\{J_\tau\}_{\tau \in \mathbb{R}}$ such that $J_\tau \in \mathcal{J}(M, \omega)$, $J_\tau \equiv J_0$ for $\tau \leq 0$ and $J_\tau \equiv J_1$ for $\tau \geq 1$. One can similarly define the spaces $\mathcal{J}^\tau(M, \omega; J_0, J_1)$ and $\mathcal{J}(M; J_0, J_1)$ of ω -tame or general 1-parameter families respectively, and all of these spaces have natural C^∞ -topologies. Given $\{J_\tau\} \in \mathcal{J}(M, \omega; J_0, J_1)$, we define the “parametrized” moduli space,

$$\mathcal{M}(\{J_\tau\}) = \{(u, \tau) \mid \tau \in \mathbb{R}, u \in \mathcal{M}(J_\tau)\},$$

along with the corresponding space of somewhere injective curves $\mathcal{M}^*(\{J_\tau\})$ and the components $\mathcal{M}_{g,m}^A(\{J_\tau\})$ for each $g, m \geq 0$, $A \in H_2(M)$. These also have natural topologies, and intuitively, we expect $\mathcal{M}^*(\{J_\tau\})$ to be a manifold whose dimension near (u, τ) is one greater than the dimension of $\mathcal{M}^*(J_\tau)$ near u . The only question is what should be the proper notion of “genericity” to make this statement correct. Given a homotopy $\{J_\tau\} \in \mathcal{J}(M, \omega; J_0, J_1)$ where $J_0, J_1 \in \mathcal{J}_{\text{reg}}(M, \omega)$, it would be too much to hope that one can always perturb $\{J_\tau\}$ so that $J_\tau \in \mathcal{J}_{\text{reg}}(M, \omega)$ for every τ ; by analogy with the case of smooth Morse functions on a manifold, any two Morse functions are indeed smoothly homotopic, but not through a family of Morse functions. What is true however is that one can find “generic homotopies,” for which $J_\tau \notin \mathcal{J}_{\text{reg}}(M, \omega)$ only for a discrete set of $\tau \in (0, 1)$, and $\mathcal{M}^*(\{J_\tau\})$ is indeed a manifold. We will not prove the following result explicitly, but the proof is an easy exercise after the proof of Theorem 4.7 is understood. Corresponding statements also hold for ω -tame or general almost complex structures.

Theorem 4.10. *Assume (M, ω) is closed and $J_0, J_1 \in \mathcal{J}_{\text{reg}}(M, \omega)$. Then there exists a Baire set $\mathcal{J}_{\text{reg}}(M, \omega; J_0, J_1) \subset \mathcal{J}(M, \omega; J_0, J_1)$ such that for every $\{J_\tau\} \in \mathcal{J}_{\text{reg}}(M, \omega; J_0, J_1)$, the parametrized space of simple curves $\mathcal{M}^*(\{J_\tau\})$ admits the structure of a smooth finite dimensional manifold, whose dimension near any $(u, \tau) \in \mathcal{M}^*(\{J_\tau\})$ is $\text{ind}(u) + 1$. Moreover, for each $\tau \in \mathbb{R}$ at which $J_\tau \in \mathcal{J}_{\text{reg}}(M, \omega)$, τ is a regular value of the natural projection $\mathcal{M}^*(\{J_\tau\}) \rightarrow \mathbb{R} : (u, \tau) \mapsto \tau$.*

The intuition behind Theorems 4.7 and 4.10 is roughly as follows. As we've already seen, spaces of J -holomorphic curves typically can be described, at least locally, as zero sets of sections of certain Banach space bundles, and we'll show in §4.3 precisely how to set up the appropriate section

$$\bar{\partial}_J : \mathcal{B} \rightarrow \mathcal{E}$$

whose zero set locally describes $\mathcal{M}_{g,m}^A(J)$. The identification between $\bar{\partial}_J^{-1}(0)$ and $\mathcal{M}_{g,m}^A(J)$ near a given curve $u \in \mathcal{M}_{g,m}^A(J)$ will in general be locally k -to-1, where k is the order of the automorphism group $\text{Aut}(u)$, and this means that even if $\bar{\partial}_J^{-1}(0)$ is a manifold, $\mathcal{M}_{g,m}^A(J)$ is at best an orbifold. This is a moot point of course if u is simple, since it then has a trivial automorphism group by Exercise 4.5. Thus once the section $\bar{\partial}_J$ is set up, the main task is to show that generic choices of J make $\bar{\partial}_J^{-1}(0)$ a manifold (at least near the simple curves), which means showing that the linearization of $\bar{\partial}_J$ is always surjective. This is a question of *transversality*, i.e. if we regard $\bar{\partial}_J$ as an embedding of \mathcal{B} into the total space \mathcal{E} and denote the zero section by $\mathcal{Z} \subset \mathcal{E}$, then $\bar{\partial}_J^{-1}(0)$ is precisely the intersection,

$$\bar{\partial}_J(\mathcal{B}) \cap \mathcal{Z},$$

and it will be a manifold if this intersection is everywhere transverse. Intuitively, one expects this to be true after a generic perturbation of $\bar{\partial}_J$, and it remains to check whether the most geometrically natural perturbation, defined by perturbing J , is “sufficiently generic” to achieve this.

The answer is yes and no: it turns out that perturbations of J are sufficiently generic if we only consider simple curves, but not for multiple covers. It's not hard to see why transversality *must* sometimes fail: if \tilde{u} is a multiple cover of u , then even if $\mathcal{M}_{\tilde{u}}(J)$ happens to be a manifold, there are certain obvious relations between the spaces $\mathcal{M}_u(J)$ and $\mathcal{M}_{\tilde{u}}(J)$ that will often cause $\mathcal{M}_{\tilde{u}}(J)$ to have “the wrong” dimension, i.e. something other than $\text{ind}(\tilde{u})$. For example, suppose $n = 4$, so M is 8-dimensional, and for some $J \in \mathcal{J}_{\text{reg}}(M, \omega)$ there exists a simple J -holomorphic sphere $u \in \mathcal{M}_{0,0}^A(J)$ with $c_1(A) = -1$. Then by (4.1), $\text{ind}(u) = 0$, and Theorem 4.7 implies that $\mathcal{M}_u^*(J)$ is a smooth 0-dimensional manifold, i.e. a discrete set. In fact the implicit function theorem implies much more (cf. Theorem 4.39): it implies that for any other $J_\epsilon \in \mathcal{J}(M, \omega)$ sufficiently close to J , there is a unique J_ϵ -holomorphic curve u_ϵ that is a small perturbation of u . Now for each of these curves and some $k \in \mathbb{N}$, consider the k -fold cover

$$\tilde{u}_\epsilon : S^2 \rightarrow M : z \mapsto u_\epsilon(z^k),$$

where as usual S^2 is identified with the extended complex plane, so that $z \mapsto z^k$ defines a k -fold holomorphic branched cover $S^2 \rightarrow S^2$. We have $[\tilde{u}_\epsilon] = k[u_\epsilon] = kA$, and thus

$$\text{ind}(\tilde{u}_\epsilon) = (n - 3)\chi(S^2) + 2c_1(kA) = 2 - 2k,$$

so if $k \geq 2$ then \tilde{u}_ϵ are J_ϵ -holomorphic spheres with negative index. By construction, these cannot be “perturbed away”: they exist for all J_ϵ sufficiently close to J , which shows that $\mathcal{M}_{\tilde{u}}(J)$ can never be made into a smooth manifold of the right dimension

by perturbing J . In this situation it is not even clear if $\mathcal{M}_{\tilde{u}}(J)$ is a manifold at all—in a few lucky situations one might be able to prove this, but it is not true in general.

The failure of Theorems 4.7 and 4.10 for multiply covered J -holomorphic curves is one of the great headaches of symplectic topology, and the major reason why fully general definitions of the various invariants based on counting holomorphic curves (Gromov-Witten theory, Floer homology, Symplectic Field Theory) are often so technically difficult as to be slightly controversial. There have been many suggested approaches to the problem, most requiring the introduction of complicated new structures, e.g. virtual moduli cycles, Kuranishi structures, polyfolds. In some fortunate situations one can avoid these complications by using topological constraints to rule out the appearance of any multiple covers in the moduli space of interest—we'll see examples of this in our applications, especially in dimension four.

4.2. Classification of pointed Riemann surfaces

4.2.1. Automorphisms and Teichmüller space. In order to understand the local structure of the moduli space of J -holomorphic curves, we will first need to consider the space of pointed Riemann surfaces, which appear as domains of such curves. In particular, we will need suitable local parametrizations of $\mathcal{M}_{g,m}$ near any given complex structure on Σ . The discussion necessarily begins with the following classical result, which is proved e.g. in [FK92].

Theorem 4.11 (Uniformization theorem). *Every simply connected Riemann surface is biholomorphically equivalent to either the Riemann sphere $S^2 = \mathbb{C} \cup \{\infty\}$, the complex plane \mathbb{C} or the upper half plane $\mathbb{H} = \{\operatorname{Im} z > 0\} \subset \mathbb{C}$.*

We will always use i to denote the standard complex structure on the Riemann sphere $S^2 = \mathbb{C} \cup \{\infty\} \cong \mathbb{C}P^1$ or the plane \mathbb{C} . The pullback of i via the diffeomorphism

$$(4.2) \quad \mathbb{R} \times S^1 \rightarrow \mathbb{C} \setminus \{0\} : (s, t) \mapsto e^{2\pi(s+it)}$$

yields a natural complex structure on the cylinder $\mathbb{R} \times S^1$, which we'll also denote by i ; it satisfies $i\partial_s = \partial_t$.

The uniformization theorem implies that every Riemann surface can be presented as a quotient of either (S^2, i) , (\mathbb{C}, i) or (\mathbb{H}, i) by some freely acting discrete group of biholomorphic transformations. We will be most interested in the punctured surfaces $(\dot{\Sigma}, j)$ where (Σ, j, Θ) is a pointed Riemann surface and $\dot{\Sigma} = \Sigma \setminus \Theta$. The only surface of this form that has S^2 as its universal cover is S^2 itself. It is almost as easy to see which surfaces are covered by \mathbb{C} , as the only biholomorphic transformations on (\mathbb{C}, i) with no fixed points are the translations, so every freely acting discrete subgroup of $\operatorname{Aut}(\mathbb{C}, i)$ is either trivial, a cyclic group of translations or a lattice. The resulting quotients are, respectively, (\mathbb{C}, i) , $(\mathbb{R} \times S^1, i) \cong (\mathbb{C} \setminus \{0\}, i)$ and the unpunctured tori (T^2, j) . All other punctured Riemann surfaces have (\mathbb{H}, i) as their universal cover, and not coincidentally, these are precisely the cases in which $\chi(\Sigma \setminus \Theta) < 0$.

Proposition 4.12. *There exists on (\mathbb{H}, i) a complete Riemannian metric g_P of constant curvature -1 that defines the same conformal structure as i and has the property that all conformal transformations on (\mathbb{H}, i) are also isometries of (\mathbb{H}, g_P) .*

PROOF. We define g_P at $z = x + iy \in \mathbb{H}$ by

$$g_P = \frac{1}{y^2} g_E,$$

where g_E is the Euclidean metric. The conformal transformations on (\mathbb{H}, i) are given by fractional linear transformations

$$\begin{aligned} \text{Aut}(\mathbb{H}, i) &= \left\{ \varphi(z) = \frac{az + b}{cz + d} \mid a, b, c, d \in \mathbb{R}, \quad ad - bc = 1 \right\} / \{\pm 1\} \\ &= \text{SL}(2, \mathbb{R}) / \{\pm 1\} =: \text{PSL}(2, \mathbb{R}), \end{aligned}$$

and one can check that each of these defines an isometry with respect to g_P . One can also compute that g_P has curvature -1 , and the geodesics of g_P are precisely the lines and semicircles that meet \mathbb{R} orthogonally, parametrized so that they exist for all forward and backward time, thus g_P is complete. For more details on all of this, the book by Hummel [Hum97] is highly recommended. \square

By lifting to universal covers, this implies the following.

Corollary 4.13. *For every pointed Riemann surface (Σ, j, Θ) such that $\chi(\Sigma \setminus \Theta) < 0$, the punctured Riemann surface $(\Sigma \setminus \Theta, j)$ admits a complete Riemannian metric g_P of constant curvature -1 that defines the same conformal structure as j , and has the property that all biholomorphic transformations on $(\Sigma \setminus \Theta, j)$ are also isometries of $(\Sigma \setminus \Theta, g_P)$.*

The metric g_P in Prop. 4.12 and Cor. 4.13 is often called the **Poincaré metric**.

The above discussion illustrates a general pattern in the study of pointed Riemann surfaces: it divides naturally into the study of punctured surfaces with negative Euler characteristic and finitely many additional cases.

Definition 4.14. A pointed surface (Σ, Θ) is said to be **stable** if $\chi(\Sigma \setminus \Theta) < 0$.

Lemma 4.15. *If (Σ, j, Θ) is a pointed Riemann surface with $\chi(\Sigma \setminus \Theta) < 0$ and $\varphi \in \text{Aut}(\Sigma, j, \Theta)$ is not the identity, then φ is also not homotopic to the identity.*

PROOF. By assumption $\varphi \neq \text{Id}$, thus by a simple unique continuation argument, it has finitely many fixed points, each of which counts with positive index since φ is holomorphic. The algebraic count of fixed points is thus at least $m = \#\Theta$. But if φ is homotopic to Id , then this count must equal $\chi(\Sigma)$ by the Lefschetz fixed point theorem, contradicting the assumption $\chi(\Sigma) < \#\Theta$. \square

The lemma implies that $\text{Aut}(\Sigma, j, \Theta)$ is always a discrete group when (Σ, Θ) is stable. In fact more is true:

Proposition 4.16. *If (Σ, j, Θ) is a closed pointed Riemann surface with either genus at least 1 or $\#\Theta \geq 3$, then $\text{Aut}(\Sigma, j, \Theta)$ is compact.*

Corollary 4.17. *If (Σ, Θ) is stable then $\text{Aut}(\Sigma, j, \Theta)$ is finite.*

Prop. 4.16 follows from the more general Lemma 4.18 below, which we'll use to show that $\mathcal{M}_{g,m}$ is Hausdorff, among other things. We should note that the corollary can be strengthened considerably, for instance one can find a priori bounds on the

order of $\text{Aut}(\Sigma, j)$ in terms of the genus, cf. [SS92, Theorem 3.9.3]. For our purposes, the knowledge that $\text{Aut}(\Sigma, j, \Theta)$ is finite will be useful enough. As we'll review below, automorphism groups in the non-stable cases are not discrete and sometimes not even compact, though they are always smooth Lie groups.

It will be convenient to have an alternative (equivalent) definition of $\mathcal{M}_{g,m}$, the moduli space of Riemann surfaces. Fix any smooth oriented closed surface Σ with genus g and an ordered set of distinct points $\Theta = (z_1, \dots, z_m) \subset \Sigma$. Then $\mathcal{M}_{g,m}$ is homeomorphic to the quotient

$$\mathcal{M}(\Sigma, \Theta) := \mathcal{J}(\Sigma) / \text{Diff}_+(\Sigma, \Theta),$$

where $\mathcal{J}(\Sigma)$ is the space of smooth almost complex structures on Σ and $\text{Diff}_+(\Sigma, \Theta)$ is the space of orientation preserving diffeomorphisms $\varphi : \Sigma \rightarrow \Sigma$ such that $\varphi|_{\Theta} = \text{Id}$. Here the action of $\text{Diff}_+(\Sigma, \Theta)$ on $\mathcal{J}(\Sigma)$ is defined by the pullback,

$$\text{Diff}_+(\Sigma, \Theta) \times \mathcal{J}(\Sigma) \rightarrow \mathcal{J}(\Sigma) : (\varphi, j) \mapsto \varphi^* j.$$

Informally speaking, $\mathcal{J}(\Sigma)$ is an infinite dimensional manifold, and we expect $\mathcal{M}(\Sigma, \Theta)$ also to be a manifold if $\text{Diff}_+(\Sigma, \Theta)$ acts freely and properly. The trouble is that in general, it does not: each $j \in \mathcal{J}(\Sigma)$ is preserved by the subgroup $\text{Aut}(\Sigma, j, \Theta)$. A solution to this complication is suggested by Lemma 4.15: if we consider not the action of all of $\text{Diff}_+(\Sigma, \Theta)$ but only the subgroup

$$\text{Diff}_0(\Sigma, \Theta) = \{\varphi \in \text{Diff}_+(\Sigma, \Theta) \mid \varphi \text{ is homotopic to Id}\},$$

then at least in the stable case, the group acts freely on $\mathcal{J}(\Sigma)$. We take this as motivation to study, as something of an intermediate step, the quotient

$$\mathcal{T}(\Sigma, \Theta) := \mathcal{J}(\Sigma) / \text{Diff}_0(\Sigma, \Theta).$$

This is the **Teichmüller space** of genus g , m -pointed surfaces. It is useful mainly because its local structure is simpler than that of $\mathcal{M}(\Sigma, \Theta)$ —we'll show below that it is always a smooth finite dimensional manifold, and its dimension can be computed using the Riemann-Roch formula. The actual moduli space of Riemann surfaces can then be understood as the quotient of Teichmüller space by a discrete group:

$$\mathcal{M}(\Sigma, \Theta) = \mathcal{T}(\Sigma, \Theta) / M(\Sigma, \Theta),$$

where $M(\Sigma, \Theta)$ is the **mapping class group**,

$$M(\Sigma, \Theta) := \text{Diff}_+(\Sigma, \Theta) / \text{Diff}_0(\Sigma, \Theta).$$

Recall that a topological group G acting continuously on a topological space X is said to act **properly** if the map $G \times X \rightarrow X \times X : (g, x) \mapsto (gx, x)$ is proper: this means that for any sequences $g_n \in G$ and $x_n \in X$ such that both x_n and $g_n x_n$ converge, g_n has a convergent subsequence. This is the condition one needs in order to show that the quotient M/G is Hausdorff. Thus for the action of $\text{Diff}_+(\Sigma, \Theta)$ or $\text{Diff}_0(\Sigma, \Theta)$ on $\mathcal{J}(\Sigma)$, we need the following compactness lemma, which also implies Prop. 4.16. We'll state it for now without proof, but will later be able to prove it using a simple case of the “bubbling” arguments in the next chapter.

Lemma 4.18. *Suppose either Σ has genus at least 1 or $\#\Theta \geq 3$. If $\varphi_k \in \text{Diff}_+(\Sigma, \Theta)$ and $j_k \in \mathcal{J}(\Sigma)$ are sequences such that $j_k \rightarrow j$ and $\varphi_k^* j_k \rightarrow j'$ in the C^∞ -topology, then φ_k has a subsequence that converges in C^∞ to a diffeomorphism $\varphi \in \text{Diff}_+(\Sigma, \Theta)$ with $\varphi^* j = j'$.*

This implies that both $\text{Diff}_+(\Sigma, \Theta)$ and $\text{Diff}_0(\Sigma, \Theta)$ act properly on $\mathcal{J}(\Sigma)$, so $\mathcal{M}(\Sigma, \Theta)$ and $\mathcal{T}(\Sigma, \Theta)$ are both Hausdorff. This is also trivially true in the cases $g = 0$, $m \leq 2$, as then $\text{Diff}_0(S^2, \Theta) = \text{Diff}_+(S^2, \Theta)$ and the uniformization theorem implies that $\mathcal{M}(S^2, \Theta) = \mathcal{T}(S^2, \Theta)$ is a one point space.

We now examine the extent to which the discrete group $M(\Sigma, \Theta)$ does not act freely on $\mathcal{T}(\Sigma, \Theta)$.

Exercise 4.19. Show that for any stable pointed Riemann surface (Σ, j, Θ) , the restriction to $\text{Aut}(\Sigma, j, \Theta)$ of the natural quotient map $\text{Diff}_+(\Sigma, \Theta) \rightarrow M(\Sigma, \Theta)$ defines an isomorphism from $\text{Aut}(\Sigma, j, \Theta)$ to the stabilizer of $[j] \in \mathcal{T}(\Sigma, \Theta)$ under the action of $M(\Sigma, \Theta)$.

Combining Exercise 4.19 with Corollary 4.17 above, we see that every point in Teichmüller space has a finite isotropy group under the action of the mapping class group; we'll see below that this is also true in the non-stable cases. This gives us the best possible picture of the local structure of $\mathcal{M}_{g,m}$: it is not a manifold in general, but locally it looks like a quotient of Euclidean space by a finite group action. Hausdorff topological spaces with this kind of local structure are called **orbifolds**. The curious reader may consult the first section of [FO99] for the definition and basic properties of orbifolds, which we will not go into here, except to state the following local structure result for $\mathcal{M}_{g,m}$.

Theorem 4.20. *$\mathcal{M}_{g,m}$ is a smooth orbifold whose isotropy subgroup at $(\Sigma, j, \Theta) \in \mathcal{M}_{g,m}$ is $\text{Aut}(\Sigma, j, \Theta)$; in particular, $\mathcal{M}_{g,m}$ is a manifold in a neighborhood of any pointed Riemann surface (Σ, j, Θ) that has trivial automorphism group. Its dimension is*

$$\dim \mathcal{M}_{g,m} = \begin{cases} 6g - 6 + 2m & \text{if } 2g + m \geq 3, \\ 2 & \text{if } g = 1 \text{ and } m = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Note that the inequality $2g + m \geq 3$ is precisely the stability condition for a genus g surface with m marked points.

The main piece of hard work that needs to be done now is proving that Teichmüller space really is a smooth manifold of the correct dimension, and in fact it will be useful to have local slices in $\mathcal{J}(\Sigma)$ that can serve as charts for $\mathcal{T}(\Sigma, \Theta)$. To that end, fix a pointed Riemann surface (Σ, j, Θ) and consider the nonlinear operator

$$\bar{\partial}_j : \mathcal{B}_\Theta^{1,p} \rightarrow \mathcal{E}^{0,p} : \varphi \mapsto T\varphi + j \circ T\varphi \circ j,$$

where $p > 2$,

$$\mathcal{B}_\Theta^{1,p} = \{\varphi \in W^{1,p}(\Sigma, \Sigma) \mid \varphi|_\Theta = \text{Id}\},$$

and $\mathcal{E}^{0,p} \rightarrow W^{1,p}(\Sigma, \Sigma)$ is the Banach space bundle with fibers

$$\mathcal{E}_\varphi^{0,p} = L^p(\overline{\text{Hom}}_\mathbb{C}(T\Sigma, \varphi^* T\Sigma)).$$

The zeroes of $\bar{\partial}_j$ are the holomorphic maps from Σ to itself that fix the marked points, and in particular a neighborhood of Id in $\bar{\partial}_j^{-1}(0)$ gives a local description of $\text{Aut}(\Sigma, j, \Theta)$. We have

$$T_{\text{Id}}\mathcal{B}_{\Theta}^{1,p} = W_{\Theta}^{1,p}(T\Sigma) := \{X \in W^{1,p}(T\Sigma) \mid X(\Theta) = 0\},$$

which is a closed subspace of $W^{1,p}(T\Sigma)$ with real codimension $2m$. The linearization

$$\mathbf{D}_{(j,\Theta)} := D\bar{\partial}_j(\text{Id}) : W_{\Theta}^{1,p}(T\Sigma) \rightarrow L^p(\overline{\text{End}}_{\mathbb{C}}(T\Sigma))$$

is then the restriction to $W_{\Theta}^{1,p}(T\Sigma)$ of the natural linear Cauchy-Riemann operator defined by the holomorphic structure of $(T\Sigma, j)$. By Riemann-Roch, the latter has index $\chi(\Sigma) + 2c_1(T\Sigma) = 3\chi(\Sigma)$, thus $\mathbf{D}_{(j,\Theta)}$ has index

$$(4.3) \quad \text{ind}(\mathbf{D}_{(j,\Theta)}) = 3\chi(\Sigma) - 2m.$$

Exercise 4.21. Show that if $A : X \rightarrow Y$ is a Fredholm operator and $X_0 \subset X$ is a closed subspace of codimension N , then $A|_{X_0}$ is also Fredholm and has index $\text{ind}(A) - N$.

Proposition 4.22. *If $\chi(\Sigma \setminus \Theta) < 0$ then $\mathbf{D}_{(j,\Theta)}$ is injective.*

PROOF. By the similarity principle, any nontrivial section $X \in \ker \mathbf{D}_{(j,\Theta)}$ has finitely many zeroes, each of positive order, and there are at least m of them since $X|_{\Theta} = 0$. Thus $\chi(\Sigma) = c_1(T\Sigma) \geq m$, which contradicts the stability assumption. \square

Observe that Prop. 4.22 provides an alternative proof of the fact that $\text{Aut}(\Sigma, j, \Theta)$ is always discrete in the stable case.

The target space of $\mathbf{D}_{(j,\Theta)}$ contains $\Gamma(\overline{\text{End}}_{\mathbb{C}}(T\Sigma))$, which one can think of as the “tangent space” to $\mathcal{J}(\Sigma)$ at j . In particular, any smooth family $j_t \in \mathcal{J}(\Sigma)$ with $j_0 = j$ has

$$\partial_t j_t|_{t=0} \in \Gamma(\overline{\text{End}}_{\mathbb{C}}(T\Sigma)).$$

We shall now use $\mathbf{D}_{(j,\Theta)}$ to define a special class of smoothly parametrized families in $\mathcal{J}(\Sigma)$.

Definition 4.23. For any $j \in \mathcal{J}(\Sigma)$, a **Teichmüller slice through j** is a smooth family of almost complex structures parametrized by an injective map

$$\mathcal{O} \rightarrow \mathcal{J}(\Sigma) : \tau \mapsto j_{\tau},$$

where \mathcal{O} is a neighborhood of 0 in some finite dimensional Euclidean space, with $j_0 = j$ and the following transversality property. If $T_j\mathcal{T} \subset \Gamma(\overline{\text{End}}_{\mathbb{C}}(T\Sigma))$ denotes the vector space of all “tangent vectors” $\partial_t j_{\tau(t)}|_{t=0}$ determined by smooth paths $\tau(t) \in \mathcal{O}$ through $\tau(0) = 0$, then

$$L^p(\overline{\text{End}}_{\mathbb{C}}(T\Sigma)) = \text{im } \mathbf{D}_{(j,\Theta)} \oplus T_j\mathcal{T}.$$

We will typically denote a Teichmüller slice simply by the image

$$\mathcal{T} := \{j_{\tau} \mid \tau \in \mathcal{O}\} \subset \mathcal{J}(\Sigma),$$

and think of this as a smoothly embedded finite dimensional submanifold of $\mathcal{J}(\Sigma)$ whose tangent space at j is $T_j\mathcal{T}$. Note that the definition doesn’t depend on p ; in

fact, one would obtain an equivalent definition by regarding $\mathbf{D}_{(j,\Theta)}$ as an operator from $W_{\Theta}^{k,p}$ to $W^{k-1,p}$ for any $k \in \mathbb{N}$ and $p > 2$.

It is easy to see that Teichmüller slices always exist. Given $j \in \mathcal{J}(\Sigma)$, pick any complement of $\text{im } \mathbf{D}_{(j,\Theta)}$, i.e. a subspace $C \subset L^p(\overline{\text{End}}_{\mathbb{C}}(T\Sigma))$ of dimension $\dim \text{coker } \mathbf{D}_{(j,\Theta)}$ whose intersection with $\text{im } \mathbf{D}_{(j,\Theta)}$ is trivial. By approximation, we may assume every section in C is smooth. We can then choose a small neighborhood $\mathcal{O} \subset C$ of 0 and define the map

$$(4.4) \quad \mathcal{O} \rightarrow \mathcal{J}(\Sigma) : y \mapsto j_y = \left(\mathbb{1} + \frac{1}{2}jy \right) j \left(\mathbb{1} + \frac{1}{2}jy \right)^{-1},$$

which has the properties $j_0 = j$ and $\partial_t j_{ty}|_{t=0} = y$, thus it is injective if \mathcal{O} is sufficiently small. This family is a Teichmüller slice through j .

Let $\pi_{\Theta} : \mathcal{J}(\Sigma) \rightarrow \mathcal{T}(\Sigma, \Theta) : j \mapsto [j]$ denote the quotient projection.

Theorem 4.24. *$\mathcal{T}(\Sigma, \Theta)$ admits the structure of a smooth finite dimensional manifold, and for any (Σ, j, Θ) there are natural isomorphisms*

$$T_{\text{Id}} \text{Aut}(\Sigma, j, \Theta) = \ker \mathbf{D}_{(j,\Theta)}, \quad T_{[j]} \mathcal{T}(\Sigma, \Theta) = \text{coker } \mathbf{D}_{(j,\Theta)}.$$

In particular,

$$(4.5) \quad \dim \mathcal{T}(\Sigma, \Theta) - \dim \text{Aut}(\Sigma, j, \Theta) = -\text{ind } \mathbf{D}_{(j,\Theta)} = 6g - 6 + 2m.$$

Moreover for any Teichmüller slice $\mathcal{T} \subset \mathcal{J}(\Sigma)$ through j , the projection

$$(4.6) \quad \pi_{\Theta}|_{\mathcal{T}} : \mathcal{T} \rightarrow \mathcal{T}(\Sigma, \Theta)$$

is a local diffeomorphism near j .

We'll prove this in the next few sections separately for the non-stable and stable cases. Observe that in the stable case, $\dim \text{Aut}(\Sigma, j, \Theta) = 0$ and thus (4.5) gives $6g - 6 + 2m$ as the dimension of Teichmüller space.

It should be intuitively clear why $\ker \mathbf{D}_{(j,\Theta)}$ is the same as $T_{\text{Id}} \text{Aut}(\Sigma, j, \Theta)$, though since $\mathbf{D}_{(j,\Theta)}$ will usually not be surjective, we still have to do something—it doesn't follow immediately from the implicit function theorem. The relationship between $T_{[j]} \mathcal{T}(\Sigma, \Theta)$ and $\text{coker } \mathbf{D}_{(j,\Theta)}$ is also not difficult to understand, though here we'll have to deal with a few analytical subtleties. Intuitively, $T_{[j]} \mathcal{T}(\Sigma, \Theta)$ should be complementary to the tangent space at $j \in \mathcal{J}(\Sigma)$ to its orbit under the action of $\text{Diff}_0(\Sigma, \Theta)$. Without worrying about the analytical details for the moment, consider a smooth family of diffeomorphisms $\varphi_{\tau} \in \text{Diff}_0(\Sigma, \Theta)$ with $\varphi_0 = \text{Id}$ and

$$\partial_{\tau} \varphi_{\tau}|_{\tau=0} = X,$$

a smooth vector field that vanishes at the marked points Θ . Then choosing a symmetric complex connection on Σ and differentiating the action $(\varphi_{\tau}, j) \mapsto \varphi_{\tau}^* j$, a short computation yields

$$(4.7) \quad \begin{aligned} \left. \frac{\partial}{\partial \tau} \varphi_{\tau}^* j \right|_{\tau=0} &= \left. \frac{\partial}{\partial \tau} [(T\varphi_{\tau})^{-1} \circ j \circ T\varphi_{\tau}] \right|_{\tau=0} = -\nabla X \circ j + j \circ \nabla X \\ &= j(\nabla X + j \circ \nabla X \circ j). \end{aligned}$$

Note that ∇ can be chosen to be the natural connection in some local holomorphic coordinates, in which case the last expression in parentheses above is simply the natural linear Cauchy-Riemann operator on $T\Sigma$ with complex structure j . Since this operator is complex linear, its image is not changed by multiplication with j , and we conclude that the tangent space to the orbit is precisely the image of $\mathbf{D}_{(j,\Theta)}$, acting on smooth vector fields that vanish at the marked points.

4.2.2. Spheres with few marked points. A pointed surface (Σ, Θ) of genus g with m marked points is stable whenever $2g + m \geq 3$. The alternative includes three cases for $g = 0$, and here uniformization tells us that (S^2, j) is equivalent to (S^2, i) for every possible j . Further, one can choose a fractional linear transformation to map up to three marked points to any points of our choosing, thus $\mathcal{M}_{g,m}$ is a one point space in each of these cases. We can now easily identify the automorphism groups for each.

- $g = 0, m = 0$: $(\Sigma, j) \cong (S^2, i)$, and $\text{Aut}(S^2, i)$ is the real 6-dimensional group of fractional linear transformations,

$$\begin{aligned} \text{Aut}(S^2, i) &= \left\{ \varphi(z) = \frac{az + b}{cz + d} \mid a, b, c, d \in \mathbb{C}, \quad ad - bc = 1 \right\} / \{\pm 1\} \\ &= \text{SL}(2, \mathbb{C}) / \{\pm 1\} =: \text{PSL}(2, \mathbb{C}). \end{aligned}$$

These are also called the **Möbius transformations**.

- $g = 0, m = 1$: $(\Sigma, j, \Theta) = (S^2, i, (\infty))$ and

$$\text{Aut}(S^2, i, (\infty)) = \text{Aut}(\mathbb{C}, i) = \{\varphi(z) = az + b \mid a, b \in \mathbb{C}\},$$

a real 4-dimensional group.

- $g = 0, m = 2$: $(\Sigma, j, \Theta) = (S^2, i, (0, \infty))$ and

$$\text{Aut}(S^2, i, (0, \infty)) = \{\varphi(z) = az \mid a \in \mathbb{C}\},$$

a real 2-dimensional group. Using the biholomorphic map (4.2), one can equivalently think of this as the group of translations on the standard cylinder $(\mathbb{R} \times S^1, i)$.

Proposition 4.25. *For each $(S^2, i, \Theta) \in \mathcal{M}_{0,m}$ with $m \leq 2$, $\mathbf{D}_{(i,\Theta)}$ is surjective and $\dim \ker \mathbf{D}_{(i,\Theta)} = \dim \text{Aut}(S^2, i, \Theta)$.*

PROOF. From (4.3), $\text{ind}(\mathbf{D}_{(i,\Theta)}) = 3\chi(\Sigma) - 2m = 6 - 2m = \dim \text{Aut}(S^2, i, \Theta)$, so it will suffice to prove that $\dim \ker \mathbf{D}_{(i,\Theta)}$ is not larger than $6 - 2m$. To see this, pick $3 - m$ distinct points $\zeta_1, \dots, \zeta_{3-m} \in \Sigma \setminus \Theta$ and consider the linear map

$$\Phi : \ker \mathbf{D}_{(i,\Theta)} \rightarrow T_{\zeta_1} \Sigma \oplus \dots \oplus T_{\zeta_{3-m}} \Sigma : X \mapsto (X(\zeta_1), \dots, X(\zeta_{3-m})).$$

The right hand side is a vector space of real dimension $6 - 2m$, so the result will follow from the claim that Φ is injective. Indeed, if $\eta \in \ker \mathbf{D}_{(i,\Theta)}$ and $\Phi(\eta) = 0$, then the similarity principle implies that each zero counts positively, and the points $\zeta_1, \dots, \zeta_{3-m}$ combined with Θ imply $c_1(TS^2) \geq 3 - m + m = 3$, giving a contradiction unless $X \equiv 0$. \square

By the above proposition, the implicit function theorem defines a smooth manifold structure on $\bar{\partial}_i^{-1}(0) \subset \mathcal{B}_\Theta^{1,p}$ near Id and yields a natural isomorphism

$$T_{\text{Id}} \text{Aut}(S^2, i, \Theta) = \ker \mathbf{D}_{(i, \Theta)}.$$

Exercise 4.26. Show that for $m \geq 3$, $\text{Aut}(S^2, i, \Theta)$ is always trivial and $\mathcal{M}_{0,m}$ is a smooth manifold of real dimension $2(m-3)$.

4.2.3. The torus. The remaining item on the list of non-stable pointed surfaces is the torus with no marked points, and this is the one case where both the automorphism groups and the Teichmüller space have positive dimension. Thus we'll see that $\mathbf{D}_{(j, \Theta)}$ is neither surjective nor injective, but fortunately the torus is a simple enough manifold so that everything can be computed explicitly.

The universal cover of (T^2, j) is the complex plane, which implies that (T^2, j) is biholomorphically equivalent to $(\mathbb{C}/\Lambda, i)$ for some lattice $\Lambda \subset \mathbb{C}$. Without loss of generality, we can take $\Lambda = \mathbb{Z} + \lambda\mathbb{Z}$ for some $\lambda \in \mathbb{H}$. Then choosing a real linear map that sends 1 to itself and λ to i , we can write $T^2 = \mathbb{C}/(\mathbb{Z} + i\mathbb{Z})$ and identify

$$(\mathbb{C}/\Lambda, i) \cong (T^2, j_\lambda),$$

where j_λ is some translation invariant complex structure on \mathbb{C} that is compatible with the standard orientation. Conversely, every such translation invariant complex structure can be obtained in this way and descends to a complex structure on T^2 .

Proposition 4.27. $[j_\lambda] = [j_{\lambda'}]$ in $\mathcal{T}(T^2)$ if and only if $\lambda = \lambda'$.

PROOF. If $j_\lambda = \varphi^* j_{\lambda'}$ for some $\varphi \in \text{Diff}_0(T^2)$, then φ can be lifted to a diffeomorphism of \mathbb{C} that (after composing with a translation) fixes the lattice $\mathbb{Z} + i\mathbb{Z}$. Now composing with the linear map mentioned above, this gives rise to a biholomorphic map $\psi : \mathbb{C} \rightarrow \mathbb{C}$ such that $\psi(0) = 0$, $\psi(1) = 1$ and $\psi(\lambda) = \lambda'$. But all biholomorphic maps on \mathbb{C} have the form $\psi(z) = az + b$, and the conditions at 0 and 1 imply $b = 0$ and $a = 1$, thus $\lambda = \lambda'$. \square

This shows that $\mathcal{T}(T^2)$ is a smooth 2-manifold that can be identified naturally with the upper half plane \mathbb{H} , and the set of translation invariant complex structures

$$\mathcal{T} := \{j_\lambda \in \mathcal{J}(T^2) \mid \lambda \in \mathbb{H}\}$$

defines a global parametrization. We'll see below that it is also a Teichmüller slice in the sense of Definition 4.23.

To understand the action of $M(T^2) = \text{Diff}_+(T^2)/\text{Diff}_0(T^2)$ on $\mathcal{T}(T^2)$, note that every element of $M(T^2)$ can be represented uniquely as a matrix $A \in \text{SL}(2, \mathbb{Z})$, which is determined by its induced isomorphism on $H_1(T^2) = \mathbb{Z}^2$. Then $A^* j_\lambda$ is another translation invariant complex structure $j_{\lambda'}$ for some $\lambda' \in \mathbb{H}$, and

$$[A] \cdot [j_\lambda] = [A^* j_\lambda] = [j_{\lambda'}].$$

Thus the stabilizer of $[j_\lambda]$ under this action is the subgroup

$$G_\lambda := \{A \in \text{SL}(2, \mathbb{Z}) \mid A^* j_\lambda = j_\lambda\}.$$

This is also a subgroup of $\text{Aut}(T^2, j_\lambda)$, and a complementary (normal) subgroup is formed by the intersection $\text{Aut}(T^2, j_\lambda) \cap \text{Diff}_0(T^2)$.

Proposition 4.28. *Every $\varphi \in \text{Aut}(T^2, j_\lambda)$ that fixes $(0, 0) \in T^2$ belongs to G_λ , and every $\varphi \in \text{Aut}(T^2, j_\lambda) \cap \text{Diff}_0(T^2)$ is a translation $\varphi(z) = z + \zeta$ for some $\zeta \in T^2$.*

PROOF. The first statement follows by a repeat of the argument used in the proof of Prop. 4.27 above: if $\varphi \in \text{Aut}(T^2, j_\lambda)$ fixes $(0, 0)$, then regarding it as a diffeomorphism on \mathbb{C}/Λ , it lifts to a biholomorphic map on \mathbb{C} which must be of the form $\psi(z) = cz$ for $c \in \mathbb{C} \setminus \{0\}$, implying that φ is the projection to $T^2 = \mathbb{C}/\mathbb{Z}^2$ of a real linear map on \mathbb{C} which preserves the lattice $\mathbb{Z} + i\mathbb{Z}$, and thus $\varphi \in \text{SL}(2, \mathbb{Z})$.

The second statement follows because one can compose any $\varphi \in \text{Aut}(T^2, j_\lambda) \cap \text{Diff}_0(T^2)$ with translations until it fixes $(0, 0)$, and conclude that the composed map is in $\text{SL}(2, \mathbb{Z}) \cap \text{Diff}_0(T^2) = \{1\}$. \square

Denoting the translation subgroup by $T^2 \subset \text{Aut}(T^2, j_\lambda)$, we see now that the total automorphism group is the semidirect product

$$\text{Aut}(T^2, j_\lambda) = T^2 \rtimes G_\lambda,$$

and is thus a smooth 2-dimensional manifold.

Proposition 4.29. *For each $[j_\lambda] \in \mathcal{T}(T^2)$, G_λ is finite.*

PROOF. The claim follows from the fact that G_λ is compact, which we show as follows. Choose a new real basis (\hat{e}_1, \hat{e}_2) for $\mathbb{C} = \mathbb{R}^2$ such that \hat{e}_1 is a positive multiple of e_1 , $\hat{e}_2 = j_\lambda \hat{e}_1$ and the parallelogram spanned by \hat{e}_1 and \hat{e}_2 has area 1. Expressing any matrix $A \in G_\lambda$ in this basis, A now belongs to both $\text{GL}(1, \mathbb{C})$ and $\text{SL}(2, \mathbb{R})$, whose intersection

$$\text{GL}(1, \mathbb{C}) \cap \text{SL}(2, \mathbb{R}) = \text{U}(1)$$

is compact. \square

By this result, $\text{Aut}(T^2, j)$ is always compact, as was predicted by Prop. 4.16. Moreover, the stabilizer of any element of $\mathcal{T}(T^2)$ under the action of $M(T^2)$ is finite, so we conclude that

$$\mathcal{M}_{2,0} \cong \mathcal{T}(T^2)/M(T^2) \cong \mathbb{H}/\text{SL}(2, \mathbb{Z})$$

is a smooth 2-dimensional orbifold, and is a manifold near any $[j_\lambda]$ for which G_λ is trivial.

Exercise 4.30. Show that G_λ is trivial for all λ in an open and dense subset of \mathbb{H} .

Let us now relate the above descriptions of $\mathcal{T}(T^2)$ and $\text{Aut}(T^2, j)$ to the natural Cauchy-Riemann operator

$$\mathbf{D}_j : W^{1,p}(TT^2) \rightarrow L^p(\overline{\text{End}}_{\mathbb{C}}(TT^2))$$

on (TT^2, j) . After an appropriate diffeomorphism we can assume without loss of generality that $j = j_\lambda \in \mathcal{T}$ for some $\lambda \in \mathbb{H}$. Then identifying TT^2 with $T^2 \times \mathbb{C}$ via the natural global complex trivialization, \mathbf{D}_j is equivalent to the standard Cauchy-Riemann operator

$$\bar{\partial} = \partial_s + i\partial_t : W^{1,p}(T^2, \mathbb{C}) \rightarrow L^p(T^2, \mathbb{C}),$$

whose kernel is the real 2-dimensional space of constant functions, which is precisely $T_{\text{Id}} \text{Aut}(T^2, j_\lambda)$ since $\text{Aut}(T^2, j_\lambda)$ consists infinitesimally of translations. Meanwhile, the formal adjoint \mathbf{D}_j^* is equivalent to

$$\partial = \partial_s - i\partial_t : W^{1,p}(T^2, \mathbb{C}) \rightarrow L^p(T^2, \mathbb{C}),$$

whose kernel is again the space of constant functions, and this is precisely $T_{j_\lambda} \mathcal{T}$.

4.2.4. The stable case. Assume $2g + m \geq 3$. We've already seen that in this case $\text{Aut}(\Sigma, j, \Theta)$ is finite and $\mathbf{D}_{(j, \Theta)}$ is injective, so Theorem 4.24 now reduces to the statement that $\mathcal{T}(\Sigma, \Theta)$ is a smooth manifold whose tangent space at $[j]$ is $\text{coker } \mathbf{D}_{(j, \Theta)}$, and local charts are given by Teichmüller slices. We argued informally above that the tangent space at $j \in \mathcal{J}(\Sigma)$ to its orbit under $\text{Diff}_0(\Sigma, \Theta)$ is the image of $\mathbf{D}_{(j, \Theta)}$, which motivates the belief that $\mathcal{T}(\Sigma, \Theta)$ should locally look like a quotient of this image, i.e. the cokernel of $\mathbf{D}_{(j, \Theta)}$.

A naive attempt to make this precise might now proceed by considering Banach manifold completions of $\mathcal{J}(\Sigma)$ and $\text{Diff}_0(\Sigma, \Theta)$ and arguing that the extension of

$$\Phi : \text{Diff}_0(\Sigma, \Theta) \times \mathcal{J}(\Sigma) \rightarrow \mathcal{J}(\Sigma) : (\varphi, j) \mapsto \varphi^* j$$

to these completions defines a smooth Banach Lie group action that is free and proper, so the quotient is a manifold whose tangent space is the quotient of the relevant tangent spaces. But this approach runs into a subtle analytical complication: the partial derivative of the map Φ with respect to the first factor must have the form

$$D_1 \Phi(\text{Id}, j)X = j \mathbf{D}_{(j, \Theta)} X,$$

and if j is not smooth, then the right hand side will always be one step less smooth than j . Indeed, $\mathbf{D}_{(j, \Theta)}$ is in this case a nonsmooth Cauchy-Riemann type operator, and we can see it more clearly by redoing the computation (4.7) in smooth coordinates that are not holomorphic: this yields a local expression of the form

$$\left. \frac{\partial}{\partial \tau} \varphi_\tau^* j \right|_{\tau=0} = j(dX + j \circ dX \circ j) + dj(X).$$

Since this involves the first derivative of j , the expression for $D_1 \Phi(\text{Id}, j)X$ can never lie in the appropriate Banach space completion of $T_j \mathcal{J}(\Sigma)$, but rather in a larger Banach space that contains it. This means that Φ is not differentiable—indeed, this is another example (cf. Exercise 2.63) of a natural map between infinite dimensional spaces that can *never* be differentiable in any conventional Banach space setting. It is probably still true that one can make a precise argument out of this idea, but it would require a significantly different analytical framework than just smooth maps on Banach manifolds, e.g. one might attempt to use the category of *sc-smooth* Banach manifolds (cf. [Hof]). Another alternative, using the correspondence between conformal structures and hyperbolic metrics on stable Riemann surfaces, is explained in [Tro92].

Instead of trying to deal with global Banach Lie group actions, we will prove the theorem by constructing smooth charts directly via local Teichmüller slices $\mathcal{T} \subset \mathcal{J}(\Sigma)$. We will indeed need to enlarge $\text{Diff}_0(\Sigma, \Theta)$ and $\mathcal{J}(\Sigma)$ to Banach manifolds containing non-smooth objects, but the key observation is that since every object in

the slice \mathcal{T} is smooth by assumption, the orbit of any $j \in \mathcal{T}$ can still be understood as a smooth Banach submanifold. The following argument was explained to me by Dietmar Salamon, on a napkin.

PROOF OF THEOREM 4.24 IN THE STABLE CASE. For $k \in \mathbb{N}$ and $p > 2$, let $\mathcal{J}^{k,p}(\Sigma)$ denote the space of $W^{k,p}$ -smooth almost complex structures on Σ , and for $k \geq 2$, let

$$\mathcal{D}_{\Theta}^{k,p} \subset W_{\Theta}^{k,p}(\Sigma, \Sigma)$$

denote the open subset consisting of all $\varphi \in W_{\Theta}^{k,p}(\Sigma, \Sigma)$ which are C^1 -smooth diffeomorphisms. Choose $j_0 \in \mathcal{J}(\Sigma)$ and suppose $\mathcal{T} \subset \mathcal{J}(\Sigma)$ is a Teichmüller slice through j_0 . This implies that $T_{j_0}\mathcal{T} \subset \Gamma(\overline{\text{End}}_{\mathbb{C}}(T\Sigma))$ is complementary to the image of

$$\mathbf{D}_{(j_0, \Theta)} : W_{\Theta}^{k,p}(T\Sigma) \rightarrow W^{k-1,p}(\overline{\text{End}}_{\mathbb{C}}(T\Sigma))$$

for all $k \in \mathbb{N}$.

Since every $j \in \mathcal{T}$ is smooth, the orbit of j under the natural action of $\mathcal{D}_{\Theta}^{k+1,p}$ is in $\mathcal{J}^{k,p}(\Sigma)$; in fact the map

$$(4.8) \quad F : \mathcal{D}_{\Theta}^{k+1,p} \times \mathcal{T} \rightarrow \mathcal{J}^{k,p}(\Sigma) : (\varphi, j) \mapsto \varphi^* j$$

is smooth and has derivative

$$\begin{aligned} dF(\text{Id}, j_0) : W_{\Theta}^{k+1,p} \oplus T_{j_0}\mathcal{T} &\rightarrow W^{k,p}(\overline{\text{End}}_{\mathbb{C}}(T\Sigma)) \\ (X, y) &\mapsto j_0 \mathbf{D}_{(j_0, \Theta)} X + y. \end{aligned}$$

This map is an isomorphism, thus by the inverse function theorem, F is a smooth diffeomorphism between open neighborhoods of $(\text{Id}, j_0) \in \mathcal{D}_{\Theta}^{k+1,p} \times \mathcal{T}$ and $j_0 \in \mathcal{J}^{k,p}(\Sigma)$.

We claim now that after shrinking \mathcal{T} if necessary, the projection $\pi_{\Theta} : \mathcal{T} \rightarrow \mathcal{T}(\Sigma, \Theta)$ is a bijection onto a neighborhood of $[j_0]$. It is clearly surjective, since every $j \in \mathcal{J}(\Sigma)$ in some neighborhood of j_0 is in the image of F . To see that it is injective, we essentially use the fact that $\text{Diff}_0(\Sigma, \Theta)$ acts freely and properly on $\mathcal{J}(\Sigma)$. Indeed, we need to show that there is no pair of sequences $j_k \neq j'_k \in \mathcal{T}$ both converging to j_0 , such that $j_k = \varphi_k^* j'_k$ for some $\varphi_k \in \text{Diff}_0(\Sigma, \Theta)$. If there are such sequences, then by Lemma 4.18, φ_k also has a subsequence converging to some $\varphi \in \text{Diff}_0(\Sigma, \Theta)$ with $\varphi^* j_0 = j_0$, thus $\varphi = \text{Id}$. But then φ_k is near the identity in $\mathcal{D}_{\Theta}^{k+1,p}$ for sufficiently large k , and $F(\varphi_k, j'_k) = j_k$ implies $(\varphi_k, j'_k) = (\text{Id}, j_k)$ since F is locally invertible.

Finally, we show that the bijection induced by any other choice of slice \mathcal{T}' through j_0 to a neighborhood of $[j_0]$ in $\mathcal{T}(\Sigma, \Theta)$ yields a smooth transition map $\mathcal{T}' \rightarrow \mathcal{T} : j' \mapsto j$. Indeed, this transition map must satisfy the relation

$$(\varphi, j) = F^{-1} \circ F'(\varphi', j')$$

for any $\varphi, \varphi' \in \mathcal{D}_{\Theta}^{k+1,p}$, where $F' : \mathcal{D}_{\Theta}^{k+1,p} \times \mathcal{T}' \rightarrow \mathcal{J}^{k,p}(\Sigma)$ is the corresponding local diffeomorphism defined for \mathcal{T}' as in (4.8). Explicitly then, $j = \text{pr}_2 \circ F^{-1} \circ F'(\text{Id}, j')$, which is clearly a smooth map. \square

Exercise 4.31. Using the Banach manifold charts constructed in the above proof, show that for any $j \in \mathcal{J}(\Sigma)$ and Teichmüller slice \mathcal{T} through j , the projection

$$L^p(\overline{\text{End}}_{\mathbb{C}}(T\Sigma)) \rightarrow T_j\mathcal{T}$$

along $\text{im } \mathbf{D}_{(j,\Theta)}$ descends to an isomorphism $\text{coker } \mathbf{D}_{(j,\Theta)} \rightarrow T_{[j]}\mathcal{T}(\Sigma, \Theta)$ that is independent of all choices.

4.3. Fredholm regularity and the implicit function theorem

With the local structure of $\mathcal{M}_{g,m}$ understood, we now turn our attention back to $\mathcal{M}(J)$, the moduli space of J -holomorphic curves. It is unfortunately not true that $\mathcal{M}(J)$ is always locally a finite dimensional manifold, nor even an orbifold. We need an extra condition to guarantee this, called *Fredholm regularity*. To understand it, we must first set up the appropriate version of the implicit function theorem.

The setup will be analogous to the case of $\mathcal{M}_{g,m}$ in the following sense. In the previous section, we analyzed $\mathcal{M}_{g,m}$ by first understanding the Teichmüller space $\mathcal{T}(\Sigma, \Theta)$. The latter is a somewhat unnatural object in that its definition depends on choices (i.e. the surface Σ and marked points $\Theta \subset \Sigma$), but it has the advantage of being a smooth finite dimensional manifold. Then the moduli space $\mathcal{M}_{g,m}$ was understood as the quotient of $\mathcal{T}(\Sigma, \Theta)$ by a discrete group action with finite isotropy groups: in fact, *locally* near a given $[j] \in \mathcal{T}(\Sigma, \Theta)$, a neighborhood in $\mathcal{M}_{g,m}$ looks like a quotient of $\mathcal{T}(\Sigma, \Theta)$ by a finite group ($\text{Aut}(\Sigma, j, \Theta)$ in the stable case), which makes $\mathcal{M}_{g,m}$ an orbifold of the same dimension as $\mathcal{T}(\Sigma, \Theta)$.

In the more general setup, we will be able to identify $\mathcal{M}_{g,m}^A(J)$ locally near a curve (Σ, j, Θ, u) with a quotient of the form

$$\bar{\partial}_J^{-1}(0) / \text{Aut}(u),$$

where $\bar{\partial}_J$ is a generalization of the nonlinear Cauchy-Riemann operator that we considered in Chapter 3, using local Teichmüller slices to incorporate varying complex structures on the domain. Its zero set thus contains all J -holomorphic curves in some neighborhood of u , but it may also include seemingly distinct curves that are actually equivalent in the moduli space, thus one must still divide by an appropriate symmetry group, which locally turns out to be the finite group $\text{Aut}(u)$. Thus $\bar{\partial}_J^{-1}(0)$ in this context plays a role analogous to that of Teichmüller space in the previous section: it is a somewhat unnatural object whose local structure is nonetheless very nice. Unlike with Teichmüller space however, the nice local structure of $\bar{\partial}_J^{-1}(0)$ doesn't come without an extra assumption, as we need the linearization of $\bar{\partial}_J$ to be a surjective operator in order to apply the implicit function theorem. When this condition is satisfied, the result will be a smooth orbifold structure for $\mathcal{M}_{g,m}^A(J)$, with its dimension determined by the index of the linearization of $\bar{\partial}$. That's the general idea; we now proceed with the details.

Suppose $(\Sigma, j, \Theta, u) \in \mathcal{M}_{g,m}^A(J)$, and choose a Teichmüller slice $\mathcal{T} \subset \mathcal{J}(\Sigma)$ through j . For any $p > 2$, denote

$$\mathcal{B}^{1,p} = W^{1,p}(\Sigma, M),$$

and define a Banach space bundle $\mathcal{E}^{0,p} \rightarrow \mathcal{T} \times \mathcal{B}^{1,p}$ whose fibers are

$$\mathcal{E}_{(j,u)}^{0,p} = L^p(\overline{\text{Hom}_{\mathbb{C}}((T\Sigma, j), (u^*TM, J))}).$$

This bundle admits the smooth section

$$\bar{\partial}_J : \mathcal{T} \times \mathcal{B}^{1,p} \rightarrow \mathcal{E}^{0,p} : (j, u) \mapsto Tu + J \circ Tu \circ j,$$

whose linearization at (j, u) is

$$(4.9) \quad \begin{aligned} D\bar{\partial}_J(j, u) : T_j\mathcal{T} \oplus W^{1,p}(u^*TM) &\rightarrow L^p(\overline{\text{Hom}_{\mathbb{C}}(T\Sigma, u^*TM)}), \\ (y, \eta) &\mapsto J \circ Tu \circ y + \mathbf{D}_u\eta, \end{aligned}$$

where on the right hand side we take j to be the complex structure on the bundle $T\Sigma$.

Definition 4.32. We say that the curve $(\Sigma, j, \Theta, u) \in \mathcal{M}_{g,m}^A(J)$ is **Fredholm regular** if the linear operator $D\bar{\partial}_J(j, u)$ of (4.9) is surjective.

The following lemma implies that our definition of Fredholm regularity doesn't depend on the choice of Teichmüller slice. Observe that it is also an open condition: if $D\bar{\partial}_J(j, u)$ is surjective then it will remain surjective after small changes in j , u and J .

Lemma 4.33. *The image of $D\bar{\partial}_J(j, u)$ doesn't depend on the choice of \mathcal{T} .*

PROOF. Let $\mathbf{L} = D\bar{\partial}_J(j, u)$ as in (4.9), and note that $T_j\mathcal{T}$ is a subspace of $L^p(\overline{\text{End}_{\mathbb{C}}(T\Sigma)})$, so \mathbf{L} can be extended to

$$\begin{aligned} \bar{\mathbf{L}} : L^p(\overline{\text{End}_{\mathbb{C}}(T\Sigma)}) \oplus T_u\mathcal{B} &\rightarrow \mathcal{E}_{(j,u)}, \\ (y, \eta) &\mapsto J \circ Tu \circ y + \mathbf{D}_u\eta. \end{aligned}$$

We claim $\text{im } \bar{\mathbf{L}} = \text{im } \mathbf{L}$. Indeed, note first that if $\mathbf{D}_{(j,\Theta)}$ denotes the natural linear Cauchy-Riemann operator on $(T\Sigma, j)$ and $y = \mathbf{D}_{(j,\Theta)}X \in \text{im } \mathbf{D}_{(j,\Theta)}$ for some $X \in W_{\Theta}^{1,p}(T\Sigma)$, then

$$\bar{\mathbf{L}}(y, 0) = J \circ Tu \circ y = Tu(jy) = Tu(\mathbf{D}_{(j,\Theta)}(jX))$$

since u is J -holomorphic and $\mathbf{D}_{(j,\Theta)}$ is complex linear. Now the following relation isn't hard to show: for any smooth vector field $X \in \Gamma(T\Sigma)$ vanishing on Θ ,

$$(4.10) \quad \mathbf{D}_u(Tu(X)) = Tu(\mathbf{D}_{(j,\Theta)}X).$$

By the density of smooth sections, this extends to all $X \in W_{\Theta}^{1,p}(T\Sigma)$, and we conclude

$$\bar{\mathbf{L}}(y, 0) \in \text{im } \mathbf{D}_u$$

whenever $y \in \text{im } \mathbf{D}_{(j,\Theta)}$. Since $L^p(\overline{\text{End}_{\mathbb{C}}(T\Sigma)}) = \text{im } \mathbf{D}_{(j,\Theta)} \oplus T_j\mathcal{T}$, it follows that \mathbf{L} and $\bar{\mathbf{L}}$ have the same image. \square

Exercise 4.34. Prove the relation (4.10) for all smooth vector fields $X \in \Gamma(T\Sigma)$. (Compare the proof of Lemma 2.30.)

Since $T_j\mathcal{T}$ is finite dimensional and \mathbf{D}_u is Fredholm, $D\bar{\partial}_J(j, u)$ is also Fredholm and has index

$$\begin{aligned}
 \text{ind } D\bar{\partial}_J(j, u) &= \dim \mathcal{T}(\Sigma, \Theta) + \text{ind } \mathbf{D}_u \\
 &= \dim \text{Aut}(\Sigma, j, \Theta) - \text{ind } \mathbf{D}_{(j, \Theta)} + \text{ind } \mathbf{D}_u \\
 (4.11) \quad &= \dim \text{Aut}(\Sigma, j, \Theta) - (3\chi(\Sigma) - 2m) + (n\chi(\Sigma) + 2c_1(u^*TM)) \\
 &= \dim \text{Aut}(\Sigma, j, \Theta) + \text{ind}(u),
 \end{aligned}$$

where we've applied (4.5), the Riemann-Roch formula and the definition of $\text{ind}(u)$.

Lemma 4.35. *For every $j \in \mathcal{J}(\Sigma)$, one can choose a Teichmüller slice \mathcal{T} through j that is invariant under the action of $\text{Aut}(\Sigma, j, \Theta)$.*

PROOF. In the case $(\Sigma, \Theta) = (T^2, \emptyset)$, one can assume after a diffeomorphism that j is translation invariant, and \mathcal{T} can then be taken to be the global Teichmüller slice defined in §4.2.3, consisting of all translation invariant complex structures compatible with the orientation. In all other cases where $\mathcal{T}(\Sigma, \Theta)$ is nontrivial, (Σ, Θ) is stable, thus the group $G := \text{Aut}(\Sigma, j, \Theta)$ is finite. Using the construction of (4.4), it suffices to find a complement $C \subset L^p(\overline{\text{End}}_{\mathbb{C}}(T\Sigma))$ of $\text{im } \mathbf{D}_{(j, \Theta)}$ that is G -invariant, as one can then compute that

$$j_{\varphi^*y} = \varphi^*j_y$$

for any $\varphi \in \text{Aut}(\Sigma, j, \Theta)$. To start with, we observe that $\text{im } \mathbf{D}_{(j, \Theta)}$ itself is G -invariant, since $\varphi^*j = j$ also implies $\mathbf{D}_{(j, \Theta)}(\varphi^*X) = \varphi^*(\mathbf{D}_{(j, \Theta)}X)$ for all $X \in W_{\Theta}^{1,p}(T\Sigma)$. A G -invariant complement C can then be defined as the L^2 -orthogonal complement of $\text{im } \mathbf{D}_{(j, \Theta)}$ with respect to any G -invariant L^2 -inner product on the sections of $\overline{\text{End}}_{\mathbb{C}}(T\Sigma)$; such a complement automatically contains only smooth sections due to linear regularity for weak solutions (cf. Corollary 2.46).

Since L^2 -inner products on $\Gamma(\overline{\text{End}}_{\mathbb{C}}(T\Sigma))$ arise naturally from j -invariant Riemannian metrics on Σ , it suffices to find such a Riemannian metric g which is also G -invariant. Recall from Corollary 4.13 that $\Sigma \setminus \Theta$ admits a complete j -invariant Riemannian metric g_P of constant curvature -1 , the *Poincaré metric*, and it has the convenient property that the biholomorphic transformations on $\Sigma \setminus \Theta$ are precisely the isometries of g_P . This is not the desired metric since it does not extend over the marked points, but we can fix this as follows: by Exercise 4.36 below, each $z \in \Theta$ admits a G -invariant neighborhood \mathcal{U}_z which can be biholomorphically identified with the unit ball $B \subset \mathbb{C}$ such that G acts by rational rotations. Thus on \mathcal{U}_z , the Euclidean metric in these coordinates is also G -invariant, and we can interpolate this with g_P near each $z \in \Theta$ to define the desired G -invariant metric on Σ . \square

Exercise 4.36. Suppose (Σ, j) is a Riemann surface and G is a finite group of biholomorphic maps on (Σ, j) which all fix the point $z \in \Sigma$. Show that z has a G -invariant neighborhood \mathcal{U}_z with a biholomorphic map $\psi : (\mathcal{U}_z, j) \rightarrow (B, i)$ such that for every $\varphi \in G$, $\psi \circ \varphi \circ \psi^{-1}$ is a rational rotation.

A quick remark about the statement of the next theorem: if $(j, u) \in \bar{\partial}_J^{-1}(0)$ and $\varphi \in \text{Aut}(\Sigma, j, \Theta)$, then the natural action

$$\varphi \cdot (j, u) = (j, u \circ \varphi)$$

preserves $\bar{\partial}_J^{-1}(0)$. Linearizing this action at $\text{Id} \in \text{Aut}(\Sigma, j, \Theta)$, we obtain a natural map of the Lie algebra $\mathfrak{aut}(\Sigma, j, \Theta)$ to $\ker D\bar{\partial}_J(j, u)$ of the form

$$\mathfrak{aut}(\Sigma, j, \Theta) \rightarrow \ker D\bar{\partial}_J(j, u) : X \mapsto (0, Tu(X)),$$

and this is an inclusion if u is not constant. Thus in the following, we can regard $\mathfrak{aut}(\Sigma, j, \Theta)$ as a subspace of $\ker D\bar{\partial}_J(j, u)$.

Theorem 4.37. *The open subset*

$$\widetilde{\mathcal{M}}_{g,m}^A(J) := \{u \in \mathcal{M}_{g,m}^A(J) \mid u \text{ is Fredholm regular and not constant}\}$$

admits the structure of a smooth finite dimensional orbifold with

$$\dim \widetilde{\mathcal{M}}_{g,m}^A(J) = \text{vir-dim } \mathcal{M}_{g,m}^A(J).$$

Its isotropy group at any $(\Sigma, j, \Theta, u) \in \mathcal{M}_{g,m}^A(J)$ is isomorphic to $\text{Aut}(u)$, so in particular, it is a manifold near u if $\text{Aut}(u)$ is trivial. There is then also a natural isomorphism

$$T_u \mathcal{M}_{g,m}^A(J) = \ker D\bar{\partial}_J(j, u) / \mathfrak{aut}(\Sigma, j, \Theta).$$

PROOF. We shall prove this in the case where $2g \geq 3$ and give some hints how to adapt the argument for the non-stable cases, leaving the details as an exercise.

Suppose $(\Sigma, j_0, \Theta, u_0) \in \mathcal{M}_{g,m}^A(J)$ is Fredholm regular and \mathcal{T} is a Teichmüller slice through j_0 which is invariant under the action of $\text{Aut}(\Sigma, j_0, \Theta)$, as supplied by Lemma 4.35. Then constructing the smooth section $\bar{\partial}_J : \mathcal{T} \times \mathcal{B}^{1,p} \rightarrow \mathcal{E}^{0,p}$ as described above, the implicit function theorem gives

$$\bar{\partial}_J^{-1}(0) \subset \mathcal{T} \times \mathcal{B}^{1,p}$$

near (j_0, u_0) the structure of a smooth submanifold with dimension $\text{ind } D\bar{\partial}_J(j_0, u_0)$. The latter is equal to $\text{ind}(u_0)$ by (4.11), since $\text{Aut}(\Sigma, j_0, \Theta)$ is in this case discrete. Moreover, since $\text{Aut}(\Sigma, j_0, \Theta)$ preserves \mathcal{T} and acts by biholomorphic maps, it also acts on $\bar{\partial}_J^{-1}(0)$ by

$$\text{Aut}(\Sigma, j_0, \Theta) \times \bar{\partial}_J^{-1}(0) \rightarrow \bar{\partial}_J^{-1}(0) : (\varphi, (j, u)) \mapsto (\varphi^* j, u \circ \varphi).$$

Clearly any two pairs related by this action correspond to equivalent curves in the moduli space, and we claim in fact that the resulting map

$$(4.12) \quad \bar{\partial}_J^{-1}(0) / \text{Aut}(\Sigma, j_0, \Theta) \rightarrow \mathcal{M}_{g,m}^A(J)$$

is a local homeomorphism onto an open neighborhood of $(\Sigma, j_0, \Theta, u_0)$. The proof of this fact uses the fact that $\text{Diff}_0(\Sigma, \Theta)$ acts freely and properly on $\mathcal{J}(\Sigma)$.

Indeed, to see that (4.12) is surjective onto a neighborhood, suppose we have a sequence $(\Sigma, j_k, \Theta, u_k) \in \mathcal{M}_{g,m}^A(J)$ with $j_k \rightarrow j_0$ and $u_k \rightarrow u_0$. Then $[j_k] \rightarrow [j_0]$ in $\mathcal{T}(\Sigma, \Theta)$, so for sufficiently large k there are unique diffeomorphisms $\varphi_k \in \text{Diff}_0(\Sigma, \Theta)$ such that $\varphi_k^* j_k$ is a sequence in \mathcal{T} approaching j_0 . Now by the properness of the action (Lemma 4.18), a subsequence of φ_k converges to an element of $\text{Aut}(\Sigma, j_0, \Theta)$ which is homotopic to the identity, and therefore is the identity since the action is also free (Lemma 4.15). It follows that $\varphi_k \rightarrow \text{Id}$, thus $u_k \circ \varphi_k \rightarrow u_0$ and for large k , $(\varphi_k^* j_k, u_k \circ \varphi_k)$ lies in an arbitrarily small neighborhood of (j_0, u_0) in $\bar{\partial}_J^{-1}(0)$.

We show now that (4.12) is injective on a sufficiently small neighborhood of (j_0, u_0) . From Exercise 4.19, $\text{Aut}(\Sigma, j_0, \Theta)$ is the stabilizer of $[j_0]$ under the action of $M(\Sigma, \Theta)$ on $\mathcal{T}(\Sigma, \Theta)$, thus the natural projection

$$\mathcal{T} / \text{Aut}(\Sigma, j_0, \Theta) \rightarrow \mathcal{M}(\Sigma, \Theta) = \mathcal{J}(\Sigma) / \text{Diff}_+(\Sigma, \Theta)$$

is a local homeomorphism near $[j_0]$. Then for any two elements (j, u) and (j', u') of $\bar{\partial}_J^{-1}(0)$ sufficiently close to (j_0, u_0) that define equivalent holomorphic curves, $[j] = [j'] \in \mathcal{M}(\Sigma, \Theta)$ implies that j and j' are related by the action of $\text{Aut}(\Sigma, j_0, \Theta)$, and this proves the claim.

We've shown that in a neighborhood of any regular $(\Sigma, j_0, \Theta, u_0) \in \mathcal{M}_{g,m}^A(J)$, the moduli space admits an orbifold chart of the correct dimension. Its isotropy group at this point is the stabilizer of (j_0, u_0) under the action of $\text{Aut}(\Sigma, j_0, \Theta)$ on $\bar{\partial}_J^{-1}(0)$, and this is precisely $\text{Aut}(u_0)$. In particular, $\mathcal{M}_{g,m}^A(J)$ is a manifold near u_0 if $\text{Aut}(u_0)$ is trivial, and the implicit function theorem identifies its tangent space at this point with $\ker D\bar{\partial}_J(j_0, u_0)$.

It remains to show that the transition maps resulting from this construction are smooth: the zero sets $\bar{\partial}_J^{-1}(0)$ inherit natural smooth structures as submanifolds of $\mathcal{T} \times \mathcal{B}^{1,p}$, but we don't yet know that these smooth structures are independent of all choices. Put another way, we need to show that for any two equivalent curves $(\Sigma, j_0, \Theta, u_0)$ and $(\Sigma', j'_0, \Theta', u'_0)$ with corresponding Teichmüller slices $\mathcal{T}, \mathcal{T}'$ and zero sets $\bar{\partial}_J^{-1}(0), (\bar{\partial}_{J'}^{-1}(0))$, there is a smooth local diffeomorphism

$$\bar{\partial}_J^{-1}(0) \rightarrow (\bar{\partial}_{J'}^{-1}(0))$$

that maps $(j_0, u_0) \mapsto (j'_0, u'_0)$ and maps each $(j, u) \in \bar{\partial}_J^{-1}(0)$ smoothly to an equivalent curve $(j', u') \in (\bar{\partial}_{J'}^{-1}(0))$. Let us just consider the case where $j_0 = j'_0$ and $u_0 = u'_0$ but the Teichmüller slices differ, as the rest is an easy exercise. For this, we can make use of the work we already did in constructing the smooth structure of Teichmüller space: if \mathcal{T} and \mathcal{T}' are two Teichmüller slices through j_0 , then there is a diffeomorphism

$$\mathcal{T} \rightarrow \mathcal{T}' : j \mapsto j'$$

such that $j' = \varphi_j^* j$ for some $\varphi_j \in \text{Diff}_0(\Sigma, \Theta)$. In fact, the diffeomorphism φ_j depends smoothly on j , as we already found a formula for it in the proof of Theorem 4.24:

$$(\varphi_j, j) = F^{-1} \circ F'(\text{Id}, j'),$$

where

$$\begin{aligned} F : \mathcal{D}_{\Theta}^{1,p} \times \mathcal{T} &\rightarrow \mathcal{J}^{0,p}(\Sigma) : (\varphi, j) \mapsto \varphi^* j, \\ F' : \mathcal{D}_{\Theta}^{1,p} \times \mathcal{T}' &\rightarrow \mathcal{J}^{0,p}(\Sigma) : (\varphi, j) \mapsto \varphi^* j \end{aligned}$$

are both smooth local diffeomorphisms near (Id, j_0) . From this formula it is clear that $\mathcal{T}' \rightarrow \mathcal{D}_{\Theta}^{1,p} : j' \mapsto \varphi_j$ is a smooth map, thus in light of the diffeomorphism between \mathcal{T} and \mathcal{T}' , so is $\mathcal{T} \rightarrow \mathcal{D}_{\Theta}^{1,p} : j \mapsto \varphi_j$. Moreover, since each φ_j is a holomorphic map $(\Sigma, j') \rightarrow (\Sigma, j)$ with both j and j' smooth, elliptic regularity implies that φ_j is also smooth. We can now define a map

$$\bar{\partial}^{-1}(0) \rightarrow \mathcal{J}^{0,p}(\Sigma) \times \mathcal{B}^{1,p} : (j, u) \mapsto (\varphi_j^* j, u \circ \varphi_j),$$

whose image is clearly in $(\bar{\partial}_J)^{-1}(0)$ and thus consists only of smooth pairs (j', u') which are equivalent to (j, u) in the moduli space. Moreover, this map is smooth since u is always smooth, again by elliptic regularity. This is the desired local diffeomorphism.

The proof is now complete for the case where (Σ, Θ) is stable. Non-stable cases come in two flavors: the simpler one is the case $g = 0$, for then Teichmüller space is trivial and we can fix $j = i$ on S^2 . Several details then simplify, except that now $\text{Aut}(S^2, i, \Theta)$ has positive dimension—nonetheless it is straightforward to see that (4.12) is still a local homeomorphism, so the only real difference in the end is the computation of the dimension,

$$\dim \mathcal{M}_{g,m}^A(J) = \text{ind } D\bar{\partial}_J(j, u) - \dim \text{Aut}(\Sigma, j, \Theta) = \text{ind}(u),$$

due to (4.11). In the case of $\mathcal{M}_{1,0}^A(J)$, for which both Teichmüller space and the automorphism groups have positive dimension, we can use the specific global Teichmüller slice of §4.2.3, and combine ideas from the stable and genus 0 cases to obtain the same result and same dimension formula in general. \square

Exercise 4.38. Work out the details of the proof of Theorem 4.37 in the non-stable cases. (*For a more detailed exposition of this in a more general context, see [Wen10a, §3.2], the proof of Theorem 0.*)

The implicit function theorem gives more than just a manifold or orbifold structure for $\mathcal{M}_{g,m}^A(J)$: it can also be used for perturbation arguments, in which the existence of curves in $\mathcal{M}_{g,m}^A(J)$ gives rise to curves in $\mathcal{M}_{g,m}^A(J')$ as well, for any J' sufficiently close to J . We stated one result along these lines already, Theorem 4.10. For another example, assume

$$\{J_\tau\}_{\tau \in (-1,1)}$$

is a smooth 1-parameter family of almost complex structures on M .

Theorem 4.39. *Suppose $(\Sigma, j_0, \Theta, u_0) \in \mathcal{M}_{g,m}^A(J_0)$ is simple and Fredholm regular with $\text{ind}(u_0) = 0$. Then for sufficiently small $\epsilon > 0$, j_0 and u_0 extend to a smooth family of complex structures j_τ and maps $u_\tau : \Sigma \rightarrow M$ for $\tau \in (-\epsilon, \epsilon)$ such that*

$$(\Sigma, j_\tau, \Theta, u_\tau) \in \mathcal{M}_{g,m}^A(J_\tau).$$

Moreover this family is unique, in the sense that for any sequence $\tau_k \rightarrow 0$ and $(\Sigma, j'_k, \Theta, u'_k) \in \mathcal{M}_{g,m}^A(J_{\tau_k})$ with $j'_k \rightarrow j_0$ and $u'_k \rightarrow u_0$ in the C^∞ -topology, we have

$$(\Sigma, j'_k, \Theta, u'_k) \sim (\Sigma, j_{\tau_k}, \Theta, u_{\tau_k})$$

for sufficiently large k .

The actual meaning of this theorem is that the *parametrized* moduli space of holomorphic curves,

$$\mathcal{M}_{g,m}^A(\{J_\tau\}) = \{(u, \tau) \mid \tau \in (-1, 1), u \in \mathcal{M}_{g,m}^A(J_\tau)\}$$

is a smooth 1-dimensional manifold near $(u_0, 0)$, and the latter is also a regular point of the projection $\mathcal{M}_{g,m}^A(\{J_\tau\}) \rightarrow (-1, 1) : (u, \tau) \mapsto \tau$. This is practically automatic, as one only has to generalize the previous setup a bit: given a Teichmüller slice \mathcal{T}

through j_0 , redefine $\mathcal{E}^{0,p}$ to be a Banach space bundle over $\mathcal{T} \times \mathcal{B}^{1,p} \times (-1, 1)$, with fiber

$$\mathcal{E}_{(j,u,\tau)}^{0,p} = L^p(\overline{\text{Hom}}_{\mathbb{C}}((T\Sigma, j), (u^*TM, J_\tau))).$$

Then we again have a natural smooth section

$$\bar{\partial} : \mathcal{T} \times \mathcal{B}^{1,p} \times (-1, 1) \rightarrow \mathcal{E}^{0,p} : (j, u, \tau) \mapsto Tu + J_\tau \circ Tu \circ j,$$

whose zero set near $(j_0, u_0, 0)$ can—at least in the absence of automorphisms—be identified with $\mathcal{M}_{g,m}^A(\{J_\tau\})$. Thus one only has to understand the linearization of $\bar{\partial}$ at $(j_0, u_0, 0)$, which is essentially the usual $D\bar{\partial}_{J_0}(j_0, u_0)$ with one extra dimension attached to the domain, raising its index to 1. If u_0 is regular then $D\bar{\partial}_{J_0}(j_0, u_0)$ is surjective and it follows immediately that $D\bar{\partial}(j_0, u_0, 0)$ is surjective as well, so the result follows as usual from the implicit function theorem.

A result of this kind can be stated more generally for any Fredholm regular curve with nonnegative index, and for this reason, regular curves are also often referred to as **unobstructed**.

4.4. Transversality for generic J

We now prove Theorem 4.7. We assume throughout this section that (M, ω) is a closed symplectic manifold and focus on ω -compatible almost complex structures, though all of our results have easily derived analogues for ω -tame or general almost complex structures (cf. Remark 4.8).

Exercise 4.40. Suppose ω is a nondegenerate 2-form on a $2n$ -dimensional vector space V , and $\mathcal{J}(V, \omega)$ denotes the space of all complex structures J on V such that $\omega(\cdot, J\cdot)$ defines a symmetric inner product. Show that $\mathcal{J}(V, \omega)$ is a smooth submanifold of $\mathcal{J}(V)$, whose tangent space at $J \in \mathcal{J}(V, \omega)$ is

$$\overline{\text{End}}_{\mathbb{C}}(V, J, \omega) := \{Y \in \overline{\text{End}}_{\mathbb{C}}(V, J) \mid \omega(v, Yw) + \omega(Yv, w) = 0 \text{ for all } v, w \in V\}.$$

Show also that for any $J_0 \in \mathcal{J}(V, \omega)$, the correspondence

$$(4.13) \quad Y \mapsto \left(\mathbb{1} + \frac{1}{2}J_0Y \right) J_0 \left(\mathbb{1} + \frac{1}{2}J_0Y \right)^{-1}$$

maps a neighborhood of 0 in $\overline{\text{End}}_{\mathbb{C}}(V, J_0, \omega)$ diffeomorphically to a neighborhood of J_0 in $\mathcal{J}(V, \omega)$.

Recall that in Exercise 4.5, we mentioned the notion of a **somewhere injective** J -holomorphic curve: this means a curve $u : \Sigma \rightarrow M$ for which the set of **injective points**,

$$\{z \in \Sigma \setminus \text{Crit}(u) \mid u^{-1}(u(z)) = \{z\}\}$$

is not empty. Somewhere injective curves are clearly simple, and the converse is also true by Prop. 4.3. As we saw in the discussion at the end of §4.1, one cannot generally expect multiply covered curves to be Fredholm regular, but it turns out

that transversality will hold generically for simple curves, precisely because they are somewhere injective. Denote

$$\mathcal{J}_{\text{reg}}(M, \omega) = \{J \in \mathcal{J}(M, \omega) \mid \text{every somewhere injective curve in } \mathcal{M}(J) \text{ is Fredholm regular}\}.$$

Theorem 4.7 can now be restated as follows.

Theorem 4.41. $\mathcal{J}_{\text{reg}}(M, \omega)$ is a Baire subset of $\mathcal{J}(M, \omega)$.

Remark 4.42. There are many possible generalizations of Theorem 4.41 to the case where M is not closed, or where $\mathcal{J}(M, \omega)$ is replaced by a more restrictive space of almost complex structures. The important observation, which will be obvious from the proof, is that if we're only interested in a set of holomorphic curves that are known to have injective points mapping to some specific compact subset $\mathcal{U} \subset M$, then it suffices to permit generic perturbations of J on \mathcal{U} but leave it fixed on $M \setminus \mathcal{U}$.

An important example that we'll see is the case of a noncompact symplectic cobordism with cylindrical ends: here M can be decomposed as

$$M = ((-\infty, 0] \times V_-) \cup M_0 \cup ([0, \infty) \times V_+),$$

where V_{\pm} are closed manifolds and M_0 is compact with $\partial M_0 = V_- \sqcup V_+$. One can then restrict attention to the space of almost complex structures that are *fixed on the cylindrical ends*, but can vary on the compact subset M_0 : a generic subset of this space then makes all simple holomorphic curves in M regular, *provided* that they also pass somewhere through M_0 . For curves that live entirely in the cylindrical ends, one can exploit the fact that V_{\pm} is compact and argue separately that a generic choice of \mathbb{R} -invariant almost complex structure on the ends achieves transversality. We will come back to this in a later chapter.

The reason to allow perturbations of J only on *compact* subsets in particular is that it is then relatively straightforward to define Banach spaces or Banach manifolds in which these perturbations live; without this assumption, one can still do something, but it requires more care.

4.4.1. Regular almost complex structures are dense. We begin by proving a weaker version of Theorem 4.41, which nonetheless suffices for most applications.

Proposition 4.43. *For each pair of integers $g, m \geq 0$ and homology class $A \in H_2(M)$, there exists a dense subset $\mathcal{J}_{\text{reg}}(g, m, A) \subset \mathcal{J}(M, \omega)$ such that for $J \in \mathcal{J}_{\text{reg}}(g, m, A)$, every somewhere injective curve $u \in \mathcal{M}_{g, m}^A(J)$ is Fredholm regular.*

Though certainly useful on its own, this statement is less beautiful than Theorem 4.41 and occasionally also less convenient. It will be the purpose of the next section to replace the word “dense” with “Baire”, using an essentially topological argument originally due to Taubes. Theorem 4.41 then follows immediately, as $\mathcal{J}_{\text{reg}}(M, \omega)$ contains the countable intersection of the sets $\mathcal{J}_{\text{reg}}(g, m, A)$ for all g, m and A , and is thus also a Baire set.

Let us sketch the proof of Prop. 4.43 before getting into the details. One must first choose a smooth Banach manifold of almost complex structures \mathcal{J}_{ϵ} in which to

vary J . One can then define a (large) separable Banach manifold that contains all the somewhere injective curves in all the moduli spaces $\mathcal{M}_{g,m}^A(J)$ for $J \in \mathcal{J}_\epsilon$, called the **universal moduli space**,

$$\mathcal{M}_{g,m}^A(\mathcal{J}_\epsilon) = \{(u, J) \mid J \in \mathcal{J}_\epsilon, u \in \mathcal{M}_{g,m}^A(J) \text{ somewhere injective}\}.$$

It takes a bit of care to make sure this space really is a Banach manifold: as usual, the main task will be to prove that a certain linear operator between Banach spaces is surjective, and this will depend crucially on the assumption that u is somewhere injective. It will also require the domain to be sufficiently large, in particular, \mathcal{J}_ϵ will have to contain a certain set of C_0^∞ -perturbations of a given J , and must therefore be infinite dimensional. Once the universal moduli space is understood, we have a natural projection map

$$\pi : \mathcal{M}_{g,m}^A(\mathcal{J}_\epsilon) \rightarrow \mathcal{J}_\epsilon : (u, J) \mapsto J,$$

whose preimage $\pi^{-1}(J)$ at any $J \in \mathcal{J}_\epsilon$ is precisely the set of all somewhere injective curves in $\mathcal{M}_{g,m}^A(J)$. Thus we expect $\mathcal{M}_{g,m}^A(J)$ to be a smooth manifold whenever J is a regular value of π , i.e. the derivative $d\pi(u, J)$ is surjective for all $(u, J) \in \pi^{-1}(J)$. In finite dimensions, Sard's theorem would tell us that this is true for almost every J , and in the present situation one can apply the following infinite dimensional version.

Theorem 4.44 (Sard-Smale). *Suppose X and Y are separable Banach spaces, $\mathcal{U} \subset X$ is an open set and $f : \mathcal{U} \rightarrow Y$ is a smooth map whose derivative $df(x) : X \rightarrow Y$ for every $x \in \mathcal{U}$ is Fredholm. Then the regular values of f form a Baire subset of Y .*

The theorem can be stated more generally for nonsmooth maps $f \in C^k(\mathcal{U}, Y)$ if k is sufficiently large, but we will not need this. We refer to [MS04, Appendix A.5] for the more general statement and a proof.

To apply the Sard-Smale theorem, we need to know that $d\pi(u, J)$ is a Fredholm operator; in fact, it turns out that $d\pi(u, J)$ has the same index and same kernel as the linearization (4.9) that defines Fredholm regularity. The latter will follow from a bit of abstract nonsense, and this immediately implies Prop. 4.43.

In fact, the argument implies that the set of regular almost complex structures is of second category in \mathcal{J}_ϵ , and you may at this point be wondering why that doesn't already prove Theorem 4.41. The answer is that we cannot simply choose \mathcal{J}_ϵ to be $\mathcal{J}(M, \omega)$, as the latter is not a Banach manifold. We are thus forced to choose a somewhat less natural space of varying almost complex structures, with a sufficiently different topology so that a Baire subset of \mathcal{J}_ϵ is not obviously a Baire subset of $\mathcal{J}(M, \omega)$, though we will easily see that it is dense. Extending density to "genericity" will require an extra argument, given in the next section.

We now carry out the details, starting with the definition of the Banach manifold \mathcal{J}_ϵ . There are two standard approaches to this: one of them, which is treated in [MS04, §3.2], is to use the space $\mathcal{J}^m(M, \omega)$ of C^m -smooth almost complex structures for sufficiently large $m \in \mathbb{N}$ and afterwards argue (using the ideas described in §4.4.2 below) that the intersection of all the spaces $\mathcal{J}_{\text{reg}}^m(M, \omega)$ gives a Baire subset of $\mathcal{J}(M, \omega)$. The disadvantage of this approach is that if J is not smooth, then the Cauchy-Riemann operator will also have only finitely many derivatives: indeed,

$\bar{\partial}_J u = Tu + J(u) \circ Tu \circ j$ involves the composition map

$$(4.14) \quad (u, J) \mapsto J \circ u$$

which may be differentiable but is not smooth unless J is (recall Lemma 2.69). Thus this approach forces one to consider Banach manifolds and maps with only finitely many derivatives, causing an extra headache that we'd hoped to avoid after we proved elliptic regularity in Chapter 2.

The alternative approach is to stay within the smooth context by defining \mathcal{J}_ϵ to be a Banach manifold that admits a continuous inclusion into $\mathcal{J}(M, \omega)$: indeed, if \mathcal{J}_ϵ embeds continuously into $\mathcal{J}^m(M, \omega)$ for every $m \in \mathbb{N}$ and u belongs to a Banach manifold such as $W^{k,p}(\Sigma, M)$, then Lemma 2.69 implies that (4.14) will be smooth. Until now, all examples we've seen of Banach spaces that embed continuously into C^∞ have been finite dimensional, and we would find such a space too small to assure the smoothness of the universal moduli space. A suitable infinite dimensional example was introduced by Floer [Flo88], and has become commonly known as the “Floer C_ϵ -space”.

Fix any $J_0 \in \mathcal{J}(M, \omega)$ and choose a sequence of positive real numbers $\epsilon_m \rightarrow 0$ for integers $m \geq 0$. Recall from Exercise 4.40 the bundle $\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega)$, whose smooth sections constitute what we think of as the “tangent space $T_{J_0}\mathcal{J}(M, \omega)$.” Define $C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$ to be the space of smooth sections $Y \in \Gamma(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$ for which the norm

$$\|Y\|_\epsilon := \sum_{m=0}^{\infty} \epsilon_m \|Y\|_{C^m}$$

is finite. This is a Banach space, and it has a natural continuous inclusion into the space of smooth sections $\Gamma(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$. It's not a priori clear from the definition that this space contains any nontrivial sections, though we do see that one can always theoretically enlarge the space by making the sequence ϵ_m converge to 0 faster. As it turns out, choosing ϵ_m small enough makes $C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$ into an infinite dimensional space that contains bump functions with small support and arbitrary values at any point:

Lemma 4.45. *Suppose $\beta : B^{2n} \rightarrow [0, 1]$ is a smooth function with compact support on the unit ball $B^{2n} \subset \mathbb{C}^n$ and $\beta(0) = 1$. One can choose a sequence of positive numbers $\epsilon_m \rightarrow 0$ such that for every $Y_0 \in \mathbb{C}^N$ and $r > 0$, the function*

$$Y(p) := \beta(p/r)Y_0$$

satisfies $\sum_{m=0}^{\infty} \epsilon_m \|Y\|_{C^m} < \infty$.

PROOF. Define $\epsilon_m > 0$ so that for $m \geq 1$,

$$\epsilon_m = \frac{1}{m^m \|\beta\|_{C^m}}.$$

Then

$$\sum_{m=1}^{\infty} \epsilon_m \|Y\|_{C^m} \leq \sum_{m=1}^{\infty} \frac{1}{m^m \|\beta\|_{C^m}} \frac{\|\beta\|_{C^m}}{r^m} = \sum_{m=1}^{\infty} \left(\frac{1/r}{m}\right)^m < \infty.$$

□

Exercise 4.46 (cf. [Flo88, Lemma 5.1]). Show that by choosing ϵ_m as in the lemma, one can arrange so that $C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$ is dense in $L^2(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$.

Now choose $\delta > 0$ sufficiently small so that the correspondence (4.13) defines an injective map

$$\{Y \in C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega)) \mid \|Y\|_\epsilon < \delta\} \rightarrow \mathcal{J}(M, \omega),$$

and define \mathcal{J}_ϵ to be its image. By construction, \mathcal{J}_ϵ is a Banach manifold (with one chart) that contains J_0 and embeds continuously into $\mathcal{J}(M, \omega)$.

As already sketched above, we now define the universal moduli space $\mathcal{M}_{g,m}^A(\mathcal{J}_\epsilon)$ to be the space of pairs (u, J) for which $J \in \mathcal{J}_\epsilon$ and u is a somewhere injective curve in $\mathcal{M}_{g,m}^A(J)$. Choosing any $(\Sigma, j_0, \Theta, u_0) \in \mathcal{M}_{g,m}^A(J_0)$ and a Teichmüller slice \mathcal{T} through j_0 , the local structure of $\mathcal{M}_{g,m}^A(\mathcal{J}_\epsilon)$ near (u_0, J_0) can be described via the zero set of a smooth section,

$$\bar{\partial} : \mathcal{T} \times \mathcal{B}^{1,p} \times \mathcal{J}_\epsilon \rightarrow \mathcal{E}^{0,p} : (j, u, J) \mapsto Tu + J \circ Tu \circ j,$$

where now $\mathcal{E}^{0,p}$ has been extended to a Banach space bundle over $\mathcal{T} \times \mathcal{B}^{1,p} \times \mathcal{J}_\epsilon$ with fiber

$$\mathcal{E}_{(j,u,J)}^{0,p} = L^p(\overline{\text{Hom}}_{\mathbb{C}}((T\Sigma, j), (u^*TM, J))).$$

The linearization $D\bar{\partial}(j_0, u_0, J_0) : T_{j_0}\mathcal{T} \oplus T_{u_0}\mathcal{B}^{1,p} \oplus T_{J_0}\mathcal{J}_\epsilon \rightarrow \mathcal{E}_{(j_0,u_0,J_0)}^{0,p}$ takes the form

$$(y, \eta, Y) \mapsto J_0 \circ Tu_0 \circ y + \mathbf{D}_{u_0}\eta + Y \circ Tu_0 \circ j_0.$$

The essential technical work is now contained in the following claim.

Proposition 4.47. *If u_0 is somewhere injective, then the operator*

$$\begin{aligned} \mathbf{L} : W^{1,p}(u_0^*TM) \oplus C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega)) &\rightarrow L^p(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, u_0^*TM)) \\ (\eta, Y) &\mapsto \mathbf{D}_{u_0}\eta + Y \circ Tu_0 \circ j_0 \end{aligned}$$

is surjective and has a bounded right inverse.

PROOF. If \mathbf{L} is surjective then the existence of a bounded right inverse follows easily since \mathbf{D}_{u_0} is Fredholm. Moreover, the Fredholm property of \mathbf{D}_{u_0} implies that $\text{im } \mathbf{L}$ is closed, thus choosing a suitable bundle metric to define the L^2 -product, it suffices to show that there is no nontrivial section $\alpha \in L^q(\overline{\text{Hom}}_{\mathbb{C}}(T\Sigma, u_0^*TM))$ with $\frac{1}{p} + \frac{1}{q} = 1$ such that

$$\begin{aligned} \langle \mathbf{D}_{u_0}\eta, \alpha \rangle_{L^2} &= 0 \quad \text{for all } \eta \in W^{1,p}(u_0^*TM), \text{ and} \\ \langle Y \circ Tu_0 \circ j_0, \alpha \rangle_{L^2} &= 0 \quad \text{for all } Y \in C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega)). \end{aligned}$$

If such α exists, then the first of these two equations implies it is a weak solution of the formal adjoint equation $\mathbf{D}_{u_0}^*\alpha = 0$, thus by regularity of weak solutions, it is smooth, and the similarity principle implies that its zero set cannot accumulate. The idea is now to choose $Y \in C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$ so that the second equation implies α *must* vanish on some nonempty open set, yielding a contradiction. There are two important details of our setup that make this possible:

- (1) u_0 has an injective point $z_0 \in \Sigma$,

- (2) $C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$ contains bump functions with small support and arbitrary values at $u_0(z_0)$.

Indeed, since the set of injective points is open and α has only isolated zeroes, we can assume without loss of generality that $\alpha(z_0) \neq 0$ and then choose (via Lemma 4.45 and Lemma 4.48 below) $Y \in C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$ so that $\langle Y \circ Tu_0 \circ j_0, \alpha \rangle$ is positive on a neighborhood of z_0 and vanishes outside this neighborhood. Then $\langle Y \circ Tu_0 \circ j_0, \alpha \rangle_{L^2}$ cannot be zero, and we have the desired contradiction. Observe the role that somewhere injectivity plays here: $Tu_0 \circ j_0$ is nonzero near z_0 since $du_0(z_0) \neq 0$, and since u_0 passes through z_0 only once (and the same is obviously true for points in a small neighborhood of z_0), fixing the value of Y near $u_0(z_0)$ only affects the L^2 -product near z_0 and nowhere else. This is why the same proof fails for multiply covered curves. \square

In choosing the bump function $Y \in C_\epsilon(\overline{\text{End}}_{\mathbb{C}}(TM, J_0, \omega))$ in the above proof, we implicitly made use of a simple linear algebra lemma. This is the only point in the argument where the symplectic structure makes any difference: it shrinks the space of available perturbations Y along J_0 , but the lemma below shows that this space is still large enough. Note that on any symplectic vector space (V, ω) with compatible complex structure J , one can choose a basis to identify J with i and ω with the standard structure ω_0 . The linear maps Y that anticommute with i and satisfy $\omega_0(Yv, w) + \omega_0(v, Yw)$ for all $v, w \in V$ are then precisely the *symmetric* matrices that are complex antilinear.

Lemma 4.48. *For any nonzero vectors $v, w \in \mathbb{R}^{2n}$, there exists a symmetric matrix Y that anticommutes with i and satisfies $Yv = w$.*

PROOF. We borrow the proof directly from [MS04, Lemma 3.2.2] and simply state a formula for Y :

$$Y = \frac{1}{|v|^2} (wv^T + vw^T + i(wv^T - vw^T)i) - \frac{1}{|v|^4} (\langle w, v \rangle (vv^T + i vv^T i) - \langle w, iv \rangle (i vv^T - vv^T i)),$$

where $\langle \cdot, \cdot \rangle$ denotes the standard real inner product on $\mathbb{R}^{2n} = \mathbb{C}^n$. \square

Since \mathcal{T} is finite dimensional, Prop. 4.47 implies that $D\bar{\partial}(j_0, u_0, J_0)$ is also surjective and has a bounded right inverse, so by the implicit function theorem, a neighborhood of (j_0, u_0, J_0) in $\bar{\partial}^{-1}(0)$ now carries the structure of a smooth Banach manifold. Repeating several details of the proof of Theorem 4.37 and exploiting the fact that $\text{Aut}(u)$ is always trivial when u is somewhere injective, it follows also that $\mathcal{M}_{g,m}^A(\mathcal{J}_\epsilon)$ is a smooth Banach manifold: locally, it can be identified with $\bar{\partial}^{-1}(0)$ and its tangent space at $((\Sigma, j, \Theta, u), J)$ is

$$T_{(u,J)}\mathcal{M}_{g,m}^A(\mathcal{J}_\epsilon) = \ker D\bar{\partial}(j, u, J).$$

Here we are restricting for notational simplicity to the case where (Σ, Θ) is stable; we leave the details of the non-stable cases as an exercise. The projection

$$\pi : \mathcal{M}_{g,m}^A(\mathcal{J}_\epsilon) \rightarrow \mathcal{J}_\epsilon : (u, J) \mapsto J$$

is obviously smooth, and its derivative at (u, J) is equivalent to the linear projection

$$\ker D\bar{\partial}(j, u, J) \rightarrow T_J\mathcal{J}_\epsilon : (y, \eta, Y) \mapsto Y.$$

Thus there is a natural identification of $\ker d\pi(u, J)$ with the kernel of $D\bar{\partial}_J(j, u)$, the operator that appears in the definition of Fredholm regularity. We claim that the cokernels of these two operators are also isomorphic, so both are Fredholm and have the same index. This is a special case of a general functional analytic fact:

Lemma 4.49. *Suppose $D : X \rightarrow Z$ is a Fredholm operator, $A : Y \rightarrow Z$ is another bounded linear operator and $L : X \oplus Y \rightarrow Z : (x, y) \mapsto Dx + Ay$ is surjective. Then the projection*

$$\Pi : \ker L \rightarrow Y : (x, y) \mapsto y$$

is Fredholm and there are natural isomorphisms $\ker \Pi = \ker D$ and $\operatorname{coker} \Pi = \operatorname{coker} D$.

PROOF. The isomorphism of the kernels is clear: it is just Π restricted to the space $\ker D \oplus \{0\}$. We construct an isomorphism $\operatorname{coker} \Pi \rightarrow \operatorname{coker} D$ as follows. Observe that $\operatorname{im} \Pi$ is simply the space of all $y \in Y$ such that $Ay = -Dx$ for any $x \in X$, hence $\operatorname{im} \Pi = A^{-1}(\operatorname{im} D)$, and

$$\operatorname{coker} \Pi = Y / \operatorname{im} \Pi = Y / A^{-1}(\operatorname{im} D).$$

Now it is easy to check that the map $A : Y \rightarrow Z$ descends to an isomorphism

$$A : Y / A^{-1}(\operatorname{im} D) \rightarrow \operatorname{im} A / (\operatorname{im} D \cap \operatorname{im} A),$$

and similarly, the inclusion $\operatorname{im} A \hookrightarrow Z$ descends to an injective homomorphism

$$\operatorname{im} A / (\operatorname{im} D \cap \operatorname{im} A) \rightarrow Z / \operatorname{im} D.$$

Since every $z \in Z$ can be written as $z = Dx + Ay$ by assumption, this map is also surjective. \square

We can now apply the Sard-Smale theorem and conclude that the regular values of π form a Baire subset of \mathcal{J}_ϵ , and for each J in this subset, Lemma 4.49 implies that the cokernel of $D\bar{\partial}_J(j, u)$ is always trivial, hence $(\Sigma, j, \Theta, u) \in \mathcal{M}_{g,m}^A(J)$ is Fredholm regular whenever it is somewhere injective. Since Baire sets are also dense, the set of regular values contains arbitrarily good approximations to J_0 in the C_ϵ -topology, and therefore also in the C^∞ -topology, implying that $\mathcal{J}_{\text{reg}}(g, m, A)$ is dense in $\mathcal{J}(M, \omega)$. The proof of Prop. 4.43 is thus complete.

4.4.2. Dense implies generic. As promised, we shall now improve Prop. 4.43 to the statement that $\mathcal{J}_{\text{reg}}(g, m, A)$ is not just dense but also of second category in $\mathcal{J}(M, \omega)$, and this implies Theorem 4.41. The idea of this step is originally due to Taubes, and it depends on the fact that the moduli space of somewhere injective J -holomorphic curves can always be exhausted—in a way that depends continuously on J —by a countable collection of compact subsets. Observe that the definition of convergence in $\mathcal{M}_{g,m}^A(J)$ does not depend in any essential way on J : thus one can sensibly speak of a convergent sequence of curves $u_k \in \mathcal{M}_{g,m}^A(J_k)$ where $J_k \in \mathcal{J}(M, \omega)$ are potentially different almost complex structures.

Lemma 4.50. *For every $J \in \mathcal{J}(M, \omega)$ and $c > 0$, there exists a subset*

$$\mathcal{M}_{g,m}^A(J, c) \subset \mathcal{M}_{g,m}^A(J)$$

such that the following properties are satisfied:

- *Every somewhere injective curve in $\mathcal{M}_{g,m}^A(J)$ is in $\mathcal{M}_{g,m}^A(J, c)$ for some $c > 0$.*
- *For each $c > 0$ and any sequence $J_k \rightarrow J$ in $\mathcal{J}(M, \omega)$, every sequence $u_k \in \mathcal{M}_{g,m}^A(J_k, c)$ has a subsequence converging to an element of $\mathcal{M}_{g,m}^A(J, c)$.*

Postponing the proof for a moment, we proceed by showing that $\mathcal{J}_{\text{reg}}(g, m, A)$ is a Baire set, because it is the intersection of a countable collection of subsets

$$\mathcal{J}_{\text{reg}}(g, m, A) = \bigcap_{c \in \mathbb{N}} \mathcal{J}_{\text{reg}, c}(g, m, A),$$

which are each open and dense in $\mathcal{J}(M, \omega)$. We define these by the condition that $J \in \mathcal{J}_{\text{reg}, c}(g, m, A)$ if and only if every curve $u \in \mathcal{M}_{g,m}^A(J, c)$ is Fredholm regular. This set obviously contains $\mathcal{J}_{\text{reg}}(g, m, A)$, and is therefore dense due to Prop. 4.43. To see that it is open, we argue by contradiction: suppose $J \in \mathcal{J}_{\text{reg}, c}(g, m, A)$ and $J_k \in \mathcal{J}(M, \omega) \setminus \mathcal{J}_{\text{reg}, c}(g, m, A)$ is a sequence converging to J . Then there is also a sequence $u_k \in \mathcal{M}_{g,m}^A(J_k, c)$ of curves that are not Fredholm regular, and a subsequence of u_k converges to some $u \in \mathcal{M}_{g,m}^A(J, c)$, which must be regular since $J \in \mathcal{J}_{\text{reg}, c}(g, m, A)$. But Fredholm regularity is an open condition, so we have a contradiction.

Theorem 4.41 is now established, except for the proof of Lemma 4.50. Morally, the lemma follows from an important fact that we haven't yet discussed but soon will: the moduli space $\mathcal{M}_{g,m}^A(J)$ has a natural compactification $\overline{\mathcal{M}}_{g,m}^A(J)$, the *Gromov compactification*, which is a metrizable topological space. In fact, one can define a metric on $\overline{\mathcal{M}}_{g,m}^A(J)$ which does not depend on J ; in a more general context, the details of this construction are carried out in [BEH⁺03, Appendix B]. Thus if we denote by $\overline{\mathcal{M}}_{\text{bad}}(J)$ the closed subset that consists of the union of $\overline{\mathcal{M}}_{g,m}^A(J) \setminus \mathcal{M}_{g,m}^A(J)$ with all the curves in $\mathcal{M}_{g,m}^A(J)$ that are multiply covered, one way to define $\mathcal{M}_{g,m}^A(J, c)$ would be as

$$\mathcal{M}_{g,m}^A(J, c) = \left\{ u \in \mathcal{M}_{g,m}^A(J) \mid \text{dist}(u, \overline{\mathcal{M}}_{\text{bad}}(J)) \geq \frac{1}{c} \right\}.$$

By Gromov's compactness theorem, any sequence $u_k \in \mathcal{M}_{g,m}^A(J_k)$ with $J_k \rightarrow J \in \mathcal{J}(M, \omega)$ has a subsequence converging to an element of $\overline{\mathcal{M}}_{g,m}^A(J)$, thus since $\mathcal{M}_{g,m}^A(J, c)$ is a closed subset of the compactification, the same holds for a sequence $u_k \in \mathcal{M}_{g,m}^A(J_k, c)$.

The above sketch is given only for the sake of intuition, as we don't actually need Gromov's compactness theorem to prove the lemma—in fact, the latter is true only for almost complex structures that are tamed by a symplectic form, and we don't need the symplectic structure either. The following proof does however contain most of the crucial analytical ingredients in the compactness theory of holomorphic curves.

PROOF OF LEMMA 4.50. We'll give a proof first for the case $g = 0$ and then sketch the modifications that are necessary for higher genus.

If $\Sigma = S^2$, then any pointed Riemann surface (Σ, j, Θ) is equivalent to one of the form (S^2, i, Θ) with $\Theta = (0, 1, \infty, z_1, \dots, z_{m-3})$ for

$$\mathbf{z} := (z_1, \dots, z_{m-3}) \in (S^2)^{m-3} \setminus \Delta,$$

where we define the subset $\Delta \subset (S^2)^{m-3}$ to consist of all tuples (z_1, \dots, z_{m-3}) such that either $z_i \in \{0, 1, \infty\}$ for some i or $z_i = z_j$ for some $i \neq j$. Choose metrics on S^2 , $(S^2)^{m-3}$ and M , with distance functions denoted by $\text{dist}(\cdot, \cdot)$. We define $\mathcal{M}_{0,m}^A(J, c)$ to be the set of all equivalence classes in $\mathcal{M}_{0,m}^A(J)$ which have representatives (S^2, i, Θ, u) with $\Theta = (0, 1, \infty, \mathbf{z})$ and the following properties:

- (1) (S^2, i, Θ) is “not close to degenerating,” in the sense that

$$\text{dist}(\mathbf{z}, \Delta) \geq \frac{1}{c}.$$

- (2) u is “not close to bubbling,” in the sense that $|du(z)| \leq c$ for all $z \in \Sigma$.

- (3) u is “not close to a multiple cover,” meaning there exists $z_0 \in \Sigma$ such that

$$\inf_{z \in \Sigma \setminus \{z_0\}} \frac{\text{dist}(u(z_0), u(z))}{\text{dist}(z_0, z)} \geq \frac{1}{c}.$$

Note that the map u is automatically somewhere injective due to the third condition, and clearly every somewhere injective (S^2, i, Θ, u) belongs to $\mathcal{M}_{0,m}^A(J, c)$ for sufficiently large c . Now if $J_k \rightarrow J \in \mathcal{J}(M, \omega)$ and we have a sequence $(S^2, i, \Theta_k, u_k) \in \mathcal{M}_{g,m}^A(J_k, c)$ with $\Theta_k = (0, 1, \infty, \mathbf{z}_k)$, we can take a subsequence so that $\mathbf{z}_k \rightarrow \mathbf{z} \in (S^2)^{m-3} \setminus \Delta$ with $\text{dist}(\mathbf{z}, \Delta) \geq 1/c$. Likewise, the uniform C^1 -bound on u_k implies a uniform $W^{1,p}$ -bound, so the regularity estimates of Chapter 2 (specifically Corollary 2.61) give a C^∞ -convergent subsequence $u_k \rightarrow u$, and we conclude $(S^2, i, \Theta_k, u_k) \rightarrow (S^2, i, \Theta, u) \in \mathcal{M}_{0,m}^A(J, c)$, where $\Theta := (0, 1, \infty, \mathbf{z})$.

For $g \geq 1$ one can no longer describe variations in j purely in terms of the marked points Θ , so we need a different trick to obtain compactness of a sequence j_k . This requires some knowledge of the Deligne-Mumford compactification of $\mathcal{M}_{g,m}$, which we will discuss in a later chapter; for now we simply summarize the main ideas. Choose a model pointed surface (Σ, Θ) with genus g and m marked points; if it is not stable, add enough additional marked points to create a stable pointed surface (Σ, Θ') , and let $m' = \#\Theta'$. Since $\chi(\Sigma \setminus \Theta') < 0$, for every $j \in \mathcal{J}(\Sigma)$ there is a unique complete hyperbolic metric g_j of constant curvature -1 on $\Sigma \setminus \Theta'$ that defines the same conformal structure as j . There is also a singular pair of pants decomposition, that is, we can fix $3g - 3 + m'$ distinct classes in $\pi_1(\Sigma \setminus \Theta')$ and choose the unique geodesic in each of these so that they separate $\Sigma \setminus \Theta'$ into $-\chi(\Sigma \setminus \Theta')$ surfaces with the homotopy type of a twice punctured disk. This procedure associates to each $j \in \mathcal{J}(\Sigma)$ a set of real numbers

$$\ell_1(j), \dots, \ell_{3g-3+m'}(j) > 0,$$

the lengths of the geodesics, which depend continuously on j . Now we define $\mathcal{M}_{g,m}^A(J, c)$ by the same scheme as with $\mathcal{M}_{0,m}^A(J, c)$ above, but replacing the first

condition by

$$\frac{1}{c} \leq \ell_i(j) \leq c \quad \text{for each } i = 1, \dots, 3g - 3 + m'.$$

Now any sequence $(\Sigma, j_k, \Theta, u_k) \in \mathcal{M}_{g,m}^A(J, c)$ has a subsequence for which the lengths $\ell_i(j_k)$ converge in $[1/c, c]$, implying that j_k converges in C^∞ to a complex structure j . The rest of the argument works as before. \square

CHAPTER 5

Bubbling and nonsqueezing

5.1. Gromov's nonsqueezing theorem

In the previous chapters we have developed a large part of the technical apparatus needed to study J -holomorphic curves in symplectic manifolds of arbitrary dimension. The only major component still missing is the compactness theory, which we will tackle in earnest in the next chapter. In this chapter we shall provide some extra motivation by explaining one of the first and most famous applications of this technical apparatus: Gromov's nonsqueezing theorem. The proof we shall give is essentially Gromov's original proof (see [Gro85, 0.3.A]), and it depends on a compactness result (Theorem 5.15) that is one of the simplest applications of Gromov's compactness theorem, but can also be proved without developing the compactness theory in its full generality. We will explain in §5.3 a proof of that result using the standard method known as “bubbling off” analysis, which also plays an essential role in the more general compactness theory.

Let us first recall the statement of the theorem. Throughout the following discussion, we shall use the symbol ω_0 to denote the standard symplectic form on Euclidean spaces of various dimensions, as well as on tori defined as

$$T^{2n} = \mathbb{R}^{2n} / N\mathbb{Z}^{2n}$$

for $N > 0$. Note that ω_0 descends to a symplectic form on T^{2n} since it is invariant under the action of \mathbb{Z}^{2n} on \mathbb{R}^{2n} by translations.

Theorem 5.1 (Gromov's “nonsqueezing” theorem [Gro85]). *For any $n \geq 2$, there exists a symplectic embedding of (B_r^{2n}, ω_0) into $(B_R^2 \times \mathbb{R}^{2n-2}, \omega_0)$ if and only if $r \leq R$.*

The existence of the embedding when $r \leq R$ is clear, so the hard part is to show that if an embedding

$$\iota : (B_r^{2n}, \omega_0) \hookrightarrow (B_R^2 \times \mathbb{R}^{2n-2}, \omega_0)$$

exists, then we must have $r \leq R$. We shall assume $r > R$ and argue by contradiction. Since the theory of J -holomorphic curves is generally easier to work with in *closed* manifolds, the first step is to transform this into a problem involving embeddings into closed symplectic manifolds. To that end, choose a small number $\epsilon > 0$ and an area form σ on the sphere S^2 such that

$$\int_{S^2} \sigma = \pi(R + \epsilon)^2.$$

Then there exists a symplectic embedding $(B_R^2, \omega_0) \hookrightarrow (S^2, \sigma)$, and hence also

$$(B_R^2 \times \mathbb{R}^{2n-2}, \omega_0) \hookrightarrow (S^2 \times \mathbb{R}^{2n-2}, \sigma \oplus \omega_0).$$

Composing this with ι above, we may regard ι as a symplectic embedding

$$\iota : (B_r^{2n}, \omega_0) \hookrightarrow (S^2 \times \mathbb{R}^{2n-2}, \sigma \oplus \omega_0).$$

We can assume without loss of generality that the image $\iota(B_r^{2n}) \subset S^2 \times \mathbb{R}^{2n-2}$ is bounded: indeed, this is obviously true for the image of a *closed* ball $\overline{B}_{r'}$ if $r' < r$, thus it can be made true for r by shrinking r slightly but keeping the condition $r > R$. We can then choose a number $N > 0$ sufficiently large so that $\iota(B_r^{2n}) \subset S^2 \times [-N, N]^{2n-2}$. Composing with the natural quotient projection on the second factor,

$$\mathbb{R}^{2n-2} \rightarrow T^{2n-2} := \mathbb{R}^{2n-2}/N\mathbb{Z}^{2n-2}$$

and letting ω_0 descend to a symplectic form on T^{2n-2} , this gives rise to a symplectic embedding

$$(5.1) \quad \iota : (B_r^{2n}, \omega_0) \rightarrow (S^2 \times T^{2n-2}, \sigma \oplus \omega_0).$$

Since $\pi_2(T^{2n-2}) = 0$, we now obtain a contradiction if we can prove the following.

Theorem 5.2. *Suppose (M, ω) is a closed symplectic manifold of dimension $2n-2 \geq 2$ which is aspherical, i.e. $\pi_2(M) = 0$, σ is an area form on S^2 , and there exists a symplectic embedding*

$$\iota : (B_r^{2n}, \omega_0) \hookrightarrow (S^2 \times M, \sigma \oplus \omega).$$

Then $\pi r^2 \leq \int_{S^2} \sigma$.

We will prove this as a corollary of the following two results. The first has its origins in the theory of minimal surfaces and is a special case of much more general results, though it admits an easy direct proof that we will explain in §5.2. The second will require us to apply the technical machinery developed in the previous chapters, together with the compactness arguments explained in §5.3.

Theorem 5.3 (monotonicity). *Suppose $r_0 > 0$, (Σ, j) is a Riemann surface and*

$$u : (\Sigma, j) \rightarrow (B_{r_0}^{2n}, i)$$

is a proper holomorphic map whose image contains 0. Then for every $r \in (0, r_0)$,

$$\int_{u^{-1}(\overline{B}_r^{2n})} u^* \omega_0 \geq \pi r^2.$$

Proposition 5.4. *Given the setup of Theorem 5.2, there exists a compatible almost complex structure $J \in \mathcal{J}(S^2 \times M, \sigma \oplus \omega)$ with $\iota^* J = i$ on B_r^{2n} and a J -holomorphic sphere*

$$u : S^2 \rightarrow S^2 \times M$$

with $[u] = [S^2 \times \{\}] \in H_2(S^2 \times M)$ whose image contains $\iota(0)$.*

Before discussing the proof of Proposition 5.4, let us prove the main result. To simplify notation, denote

$$(W, \Omega) := (S^2 \times M, \sigma \oplus \omega), \quad \text{and} \quad A_0 := [S^2 \times \{*\}] \in H_2(W).$$

Recall that in Chapter 2, we defined the *energy* $E(u)$ of a J -holomorphic curve $u : \Sigma \rightarrow W$ as $\int_{\Sigma} u^* \Omega$, and observed that whenever J is tamed by Ω , this is also

the (nonnegative!) area traced out by u for a natural choice of Riemannian metric on W . For the curve $u : S^2 \rightarrow S^2 \times W$ provided by Proposition 5.4, we can find the energy by a purely homological computation:

$$E(u) = \int_{S^2} u^* \Omega = \langle [\Omega], [u] \rangle = \langle [\sigma \oplus \omega], A_0 \rangle = \langle [\sigma], [S^2] \rangle = \int_{S^2} \sigma.$$

Since the integrand $u^* \Omega$ is always nonnegative, this gives an upper bound for the amount of energy u has in the image of the ball B_r^{2n} , and in this ball, we can use ι^{-1} to pull back u to a map $\iota^{-1} \circ u : u^{-1}(\iota(B_r^{2n})) \rightarrow B_r^{2n}$ which contains 0 in its image and is i -holomorphic since $\iota^* J = i$. Thus combining the above upper bound with the lower bound from Theorem 5.3, we find that for any $r' \in (0, r)$,

$$\pi(r')^2 \leq \int_{u^{-1}(\iota(B_{r'}^{2n}))} (\iota^{-1} \circ u)^* \omega_0 = \int_{u^{-1}(\iota(B_{r'}^{2n}))} u^* \Omega \leq \int_{S^2} u^* \Omega = \int_{S^2} \sigma.$$

This proves Theorem 5.2.

For the rest of this section, we discuss the truly nontrivial part of the proof above: why does the J -holomorphic sphere in Proposition 5.4 exist? This turns out to be true not just for a specific J but also for *generic* Ω -compatible almost complex structures on W , and there is nothing special about the point $\iota(0)$, as *every* point in W is in the image of some J -holomorphic sphere homologous to A_0 . Moreover, this is also true for a generic subset of the special class of almost complex structures that match the integrable complex structure $\iota_* i$ on $\iota(B_r^{2n})$. We will not be able to find these J -holomorphic curves explicitly, as we have no concrete knowledge about the symplectic embedding $\iota : B_r^{2n} \rightarrow W$ and thus cannot even write down an explicit expression for J having the desired property in $\iota(B_r^{2n})$. Instead, we argue from more abstract principles by starting from a simpler almost complex structure, for which the holomorphic curves are easy to classify, and then using a deformation argument to show that the desired curves for our more general data must also exist. This argument can be outlined as follows:

- (1) Find a special $J_0 \in \mathcal{J}(W, \Omega)$ for which the moduli space $\mathcal{M}_{0,1}^{A_0}(J_0)$ of J_0 -holomorphic spheres homologous to $[S^2 \times \{*\}]$ and with one marked point is easy to describe precisely: in particular, the curves in $\mathcal{M}_{0,1}^{A_0}(J_0)$ are all Fredholm regular, and the moduli space is a closed $2n$ -dimensional manifold diffeomorphic to W , with a diffeomorphism provided by the natural evaluation map

$$\text{ev} : \mathcal{M}_{0,1}^{A_0}(J_0) \rightarrow W : [(S^2, j, z, u)] \mapsto u(z).$$

- (2) Choose $J_1 \in \mathcal{J}(W, \Omega)$ with the desired property $\iota^* J_1 = i$ and show that for a generic such choice, the moduli space $\mathcal{M}_{0,1}^{A_0}(J_1)$ is also a smooth $2n$ -dimensional manifold.
- (3) Choose a homotopy $\{J_t\}$ from J_0 to J_1 and show that for a generic such choice, the resulting parametrized moduli space $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$ is a smooth $(2n + 1)$ -dimensional manifold with boundary

$$\partial \mathcal{M}_{0,1}^{A_0}(\{J_t\}) = \mathcal{M}_{0,1}^{A_0}(J_0) \sqcup \mathcal{M}_{0,1}^{A_0}(J_1).$$

Moreover, $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$ is *compact*.

- (4) Since $\text{ev} : \mathcal{M}_{0,1}^{A_0}(J_0) \rightarrow W$ is a diffeomorphism, its \mathbb{Z}_2 -mapping degree is 1, and the fact that ev extends naturally over the cobordism $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$ implies that its restriction to the other boundary component $\mathcal{M}_{0,1}^{A_0}(J_1)$ also has \mathbb{Z}_2 -degree 1. It follows that $\text{ev} : \mathcal{M}_{0,1}^{A_0}(J_1) \rightarrow W$ is surjective, so for every $p \in W$, there is a J_1 -holomorphic sphere $u : S^2 \rightarrow W$ with $[u] = A_0$ and a point $z \in S^2$ such that $u(z) = p$.

We carry out the details in the next several subsections. The only part that cannot be proved using the tools we've already developed is the compactness of $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$, which is incidentally the only place where the assumption $\pi_2(M) = 0$ is used. This compactness is a deep result which we shall prove in §5.3.

5.1.1. The moduli space for J_0 . Identify S^2 with the Riemann sphere $\mathbb{C} \cup \{\infty\}$ with its standard complex structure i , choose any $J_M \in \mathcal{J}(M, \omega)$, and define $J_0 \in \mathcal{J}(W, \Omega)$ via the natural direct sum decomposition $T_{(z,p)}W = T_z S^2 \oplus T_p M$, that is

$$J_0 := i \oplus J_M.$$

Then a map $u = (u_S, u_M) : S^2 \rightarrow S^2 \times M$ is J_0 -holomorphic if and only if $u_S : S^2 \rightarrow S^2$ is holomorphic and $u_M : S^2 \rightarrow M$ is J_M -holomorphic. If $[u] = A_0 = [S^2 \times \{*\}]$, then we also have

$$[u_S] = [S^2], \quad \text{and} \quad [u_M] = 0.$$

The latter implies that u_M has zero energy as a J_M -holomorphic curve in M , i.e. $\int_{S^2} u_M^* \omega = \langle [\omega], [u_M] \rangle = 0$, hence u_M is constant. Moreover, $u_S : S^2 \rightarrow S^2$ is a holomorphic map of degree 1, and thus is biholomorphic (cf. Exercise 2.90), so after a reparametrization of the domain we can assume $u_S = \text{Id}$. It follows that the moduli space $\mathcal{M}_{0,1}^{A_0}(J_0)$ can be identified with the following set:

$$\mathcal{M}_{0,1}^{A_0}(J_0) = \{(u_m, \zeta) \mid m \in M \text{ and } \zeta \in S^2\},$$

where we define the J_0 -holomorphic maps

$$u_m : S^2 \rightarrow S^2 \times M : z \mapsto (z, m).$$

The evaluation map $\text{ev} : \mathcal{M}_{0,1}^{A_0}(J_0) \rightarrow S^2 \times M$ then takes the form

$$\text{ev}(u_m, \zeta) = (\zeta, m),$$

and is thus clearly a diffeomorphism. Observe that there is a natural splitting of complex vector bundles

$$(5.2) \quad u_m^* TW = TS^2 \oplus E_0^{(n-1)},$$

where $E_0^{(n-1)} \rightarrow S^2$ denotes the trivial complex bundle of rank $n - 1$ whose fiber at every point $z \in S^2$ is $(T_z M, J_M)$.

The observations above imply that $\mathcal{M}_{0,1}^{A_0}(J_0)$ is a smooth manifold of dimension $2n$, and indeed, this is precisely the prediction made by the index formula (4.1), which gives

$$\text{vir-dim } \mathcal{M}_{0,1}^{A_0}(J_0) = 2(n - 3) + 2c_1(A_0) + 2 = 2n$$

after plugging in the computation

$$c_1(A_0) = c_1(u_m^* T(S^2 \times M)) = c_1(TS^2) + c_1(E_0^{n-1}) = 2.$$

The above does *not* immediately imply that every curve in $\mathcal{M}_{0,1}^{A_0}(J_0)$ is Fredholm regular; in general only the converse of this statement is true. This is something we will need to know in order to understand the local structure of the parametrized moduli space $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$, and we must proceed with caution since our choice of J_0 is definitively *non-generic*.¹ This means that we cannot expect transversality to be achieved for general reasons, but must instead check it explicitly. This turns out to be not so hard, simply because the curves $u_m(z) = (z, m)$ are so explicit.

Lemma 5.5. *Every J_0 -holomorphic sphere of the form $u_m : S^2 \rightarrow S^2 \times M : z \mapsto (z, m)$ for $m \in M$ is Fredholm regular.*

PROOF. We recall from Definition 4.32 that u_m is Fredholm regular if and only if a certain bounded linear operator of the form

$$D\bar{\partial}_{J_0}(i, u_m) : T_i \mathcal{T} \oplus W^{1,p}(u_m^* TW) \rightarrow L^p(\overline{\text{Hom}}_{\mathbb{C}}(TS^2, u_m^* TW))$$

is surjective. Here \mathcal{T} is a Teichmüller slice, which in the present case is trivial since the Teichmüller space of S^2 with one marked point is trivial, so we can drop this factor and simply consider the linearized Cauchy-Riemann operator

$$\mathbf{D}_{u_m} : W^{1,p}(u_m^* TW) \rightarrow L^p(\overline{\text{Hom}}_{\mathbb{C}}(TS^2, u_m^* TW)).$$

We can make use of the natural splitting (5.2) to split the domain and target of \mathbf{D}_{u_m} as

$$W^{1,p}(u_m^* TW) = W^{1,p}(TS^2) \oplus W^{1,p}(E_0^{n-1})$$

and

$$L^p(\overline{\text{Hom}}_{\mathbb{C}}(TS^2, u_m^* TW)) = L^p(\overline{\text{End}}_{\mathbb{C}}(TS^2)) \oplus L^p(\overline{\text{Hom}}_{\mathbb{C}}(TS^2, E_0^{n-1})).$$

In light of the split nature of the nonlinear Cauchy-Riemann equation for J_0 -holomorphic maps $u : S^2 \rightarrow S^2 \times M$, it then turns out that the matrix form of \mathbf{D}_{u_m} with respect to these splittings is

$$\mathbf{D}_{u_m} = \begin{pmatrix} \mathbf{D}_i^{S^2} & 0 \\ 0 & \mathbf{D}_m \end{pmatrix},$$

where $\mathbf{D}_i^{S^2} : W^{1,p}(TS^2) \rightarrow L^p(\overline{\text{End}}_{\mathbb{C}}(TS^2))$ is the natural Cauchy-Riemann operator defined by the holomorphic vector bundle structure of (TS^2, i) , and

$$\mathbf{D}_m : W^{1,p}(T_m M) \rightarrow L^p(\overline{\text{Hom}}_{\mathbb{C}}(TS^2, T_m M))$$

is the linearization of $\bar{\partial}_{J_M}$ at the constant J_M -holomorphic sphere $S^2 \rightarrow M : z \mapsto m$. Specializing (2.9) for the case of a constant map, we see that the latter is simply the standard Cauchy-Riemann operator on the trivial bundle E_0^{n-1} , i.e. it is the operator determined by the unique holomorphic structure on E_0^{n-1} for which the constant sections are holomorphic. As such, this operator splits further with respect

¹Even if $j \in \mathcal{J}(S^2)$ and $J_M \in \mathcal{J}(M, \omega)$ are chosen generically, product structures of the form $j \oplus J_M$ on $S^2 \times M$ are still of a rather special type that can never be regarded as generic. See Remark 5.7 for an example of just how badly things can potentially go wrong.

to the splitting of E_0^{n-1} into holomorphic line bundles determined by any complex basis of $T_m M$. This yields a presentation of \mathbf{D}_{u_m} in the form

$$\mathbf{D}_{u_m} = \begin{pmatrix} \mathbf{D}_i^{S^2} & 0 & \cdots & 0 \\ 0 & \bar{\partial} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \bar{\partial} \end{pmatrix},$$

where each of the diagonal terms are complex linear Cauchy-Riemann type operators on line bundles, with the $\bar{\partial}$ entries in particular denoting operators that are equivalent to the standard operator

$$\bar{\partial} : W^{1,p}(S^2, \mathbb{C}) \rightarrow L^p(\overline{\text{Hom}}_{\mathbb{C}}(TS^2, \mathbb{C})) : f \mapsto df + i df \circ i.$$

These operators are surjective by Theorem 3.20 since $c_1(E_0^1) = 0 > -\chi(S^2)$. Similarly, $\mathbf{D}_i^{S^2}$ is also surjective since $c_1(TS^2) = 2 > -\chi(S^2)$. \square

Remark 5.6. The above is an example of a general phenomenon often called “automatic transversality”: it refers to various situations in which despite (or in this case even *because of*) a non-generic choice of J , transversality can be achieved by reducing it to a problem involving Cauchy-Riemann operators on line bundles and applying Theorem 3.20. The case above is unusually fortunate, as it is not often possible to split a given Cauchy-Riemann operator over a sum of line bundles in just the right way. In dimension four, however, arguments like this do often work out in greater generality, and we’ll make considerable use of them in later applications to symplectic 4-manifolds.

Remark 5.7. The following example is meant to persuade you that no almost complex structure of the product form $j \oplus J_M$ can be regarded as “generic” by any reasonable definition. Suppose (Σ, j) is a closed connected Riemann surface of genus g , σ is a compatible area form on Σ , $J_M \in \mathcal{J}(M, \omega)$ is as above and $J_0 = j \oplus J_M \in \mathcal{J}(\Sigma \times M, \sigma \oplus \omega)$. Then any J_0 -holomorphic curve of the form $u_m : \Sigma \rightarrow \Sigma \times M : z \mapsto (z, m)$ for $m \in M$ has

$$c_1(u_m^* T(\Sigma \times M)) = c_1(T\Sigma) = \chi(\Sigma),$$

so

$$\text{ind}(u_m) = (n - 3)\chi(\Sigma) + 2c_1([\Sigma \times \{*\}]) = (n - 1)\chi(\Sigma),$$

which for $n \geq 2$ is negative whenever $g \geq 2$. Thus in this case, a generic perturbation of J_0 should eliminate such curves altogether, but it is clear that a perturbation of the form $i \oplus J'_M$ for $J'_M \in \mathcal{J}(M, \omega)$ will never accomplish this. In Lemma 5.5, we were simply lucky to be working with genus zero.

5.1.2. Transversality for J_1 . From now on, assume the symplectic embedding $\iota : (B_r^{2n}, \omega_0) \rightarrow (W, \Omega)$ can be extended symplectically to a neighborhood of the closure \overline{B}_r^{2n} ; this can always be achieved by shrinking r slightly without violating the assumption $r > R$. Now consider the closed subspace of $\mathcal{J}(W, \Omega)$ defined by

$$\mathcal{J}(W, \Omega; \iota) := \{J \in \mathcal{J}(W, \Omega) \mid \iota^* J = i \text{ on } \overline{B}_r^{2n}\},$$

in other words this is the space of all Ω -compatible almost complex structures on W which match the particular integrable complex structure ι_*i on the closed set $\iota(\overline{B}_r^{2n})$.

Exercise 5.8. Convince yourself that $\mathcal{J}(W, \Omega; \iota)$ is not empty. *Hint: it may help to recall that the usual space of compatible almost complex structures is always not only nonempty but also connected, see §2.2.*

As with J_0 in the previous subsection, the condition $\iota^*J = i$ is nongeneric in some sense, but it turns out not to matter for our purposes:

Proposition 5.9. *There exists a Baire subset $\mathcal{J}_{\text{reg}}(W, \Omega; \iota) \subset \mathcal{J}(W, \Omega; \iota)$ such that for any $J \in \mathcal{J}_{\text{reg}}(W, \Omega; \iota)$, all J -holomorphic spheres homologous to A_0 are Fredholm regular, hence $\mathcal{M}_{0,1}^{A_0}(J)$ is a smooth manifold of dimension $2n$.*

PROOF. We begin with the following observations:

- (1) The virtual dimension of $\mathcal{M}_{g,m}^A(J)$ depends in general on g , m and A , but not on J , thus our earlier computation $\text{vir-dim } \mathcal{M}_{0,1}^{A_0}(J_0) = 2n$ also applies to $\mathcal{M}_{0,1}^{A_0}(J)$ for any J .
- (2) Every pseudoholomorphic curve $u : S^2 \rightarrow W$ homologous to A_0 is *simple*, as $A_0 = [S^2 \times \{*\}]$ is not a positive multiple of any other homology class in $H_2(S^2 \times M)$.
- (3) For any $J \in \mathcal{J}(W, \Omega; \iota)$, there is no closed nonconstant J -holomorphic curve $u : \Sigma \rightarrow W$ whose image lies entirely in $\iota(\overline{B}_r^{2n})$. If such a curve did exist, then $\iota^{-1} \circ u$ would be a nonconstant closed i -holomorphic curve in \mathbb{R}^{2n} and would thus have positive energy

$$\int_{\Sigma} (\iota^{-1} \circ u)^* \omega_0 > 0,$$

but this is impossible since ω_0 vanishes on every cycle in \mathbb{R}^{2n} .

The result now follows by a minor modification of the proof of Theorem 4.7, see also Remark 4.42. The crucial point is that the set of perturbations allowed by $\mathcal{J}(W, \Omega; \iota)$ is still large enough to prove that the universal moduli space for somewhere injective curves is smooth, because every such curve necessarily has an injective point outside of $\iota(\overline{B}_r^{2n})$. \square

In light of this result, we can choose

$$J_1 \in \mathcal{J}_{\text{reg}}(W, \Omega; \iota)$$

so that $\mathcal{M}_{0,1}^{A_0}(J_1)$ is a smooth manifold of dimension $2n$.

5.1.3. The homotopy of almost complex structures. Denote by

$$\mathcal{J}(W, \Omega; J_0, J_1),$$

the space of smooth Ω -compatible homotopies between J_0 and J_1 , i.e. this consists of all smooth 1-parameter families $\{J_t\}_{t \in [0,1]}$ such that $J_t \in \mathcal{J}(W, \Omega)$ for all $t \in [0, 1]$ and J_t matches the structures chosen above for $t = 0, 1$. This gives rise to the parametrized moduli space

$$\mathcal{M}_{0,1}^{A_0}(\{J_t\}) = \{(u, t) \mid t \in [0, 1], u \in \mathcal{M}(J_t)\}.$$

The following is the fundamental input we need from the compactness theory of holomorphic curves. It depends on certain topological details in the setup we've chosen, and in particular on the fact that $A_0 = [S^2 \times \{*\}]$ is a primitive homology class and $\pi_2(M) = 0$.

Proposition 5.10. *For any $\{J_t\} \in \mathcal{J}(W, \Omega; J_0, J_1)$, $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$ is compact.*

We'll come back to the proof of this in §5.3. Notice that since $\mathcal{M}_{0,1}^{A_0}(J_1)$ is naturally a closed subset of $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$ and is already known to be a smooth manifold, this implies that $\mathcal{M}_{0,1}^{A_0}(J_1)$ is a closed manifold. Since Fredholm regularity is an open condition, the same is then true for all $\mathcal{M}_{0,1}^{A_0}(J_t)$ with t in some neighborhood of either 0 or 1, and for t in this range the natural projection

$$\mathcal{M}_{0,1}^{A_0}(\{J_t\}) \rightarrow \mathbb{R} : (u, t) \mapsto t$$

is a submersion. We cannot expect this to be true for all $t \in [0, 1]$, not even for a generic choice of the homotopy, but by applying Theorem 4.10 we can at least arrange for $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$ to carry a smooth structure:

Proposition 5.11. *There exists a Baire subset*

$$\mathcal{J}_{\text{reg}}(W, \Omega; J_0, J_1) \subset \mathcal{J}(W, \Omega; J_0, J_1)$$

such that for any $\{J_t\} \in \mathcal{J}_{\text{reg}}(W, \Omega; J_0, J_1)$, $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$ is a compact smooth manifold, with boundary

$$\partial \mathcal{M}_{0,1}^{A_0}(\{J_t\}) = \mathcal{M}_{0,1}^{A_0}(J_0) \sqcup \mathcal{M}_{0,1}^{A_0}(J_1).$$

5.1.4. Conclusion of the proof. We will now derive the desired existence result using the \mathbb{Z}_2 -mapping degree of the evaluation map. Recall that in general, if X and Y are closed and connected n -dimensional manifolds and $f : X \rightarrow Y$ is a continuous map, then the degree $\deg_2(f) \in \mathbb{Z}_2$ can be defined by the condition

$$f_*[X] = \deg_2(f)[Y] \in H_n(Y; \mathbb{Z}_2),$$

where $[X] \in H_n(X; \mathbb{Z}_2)$ and $[Y] \in H_n(Y; \mathbb{Z}_2)$ denote the respective fundamental classes with \mathbb{Z}_2 -coefficients. Equivalently, if f is smooth then $\deg_2(f)$ can be defined as the modulo 2 count of points in $f^{-1}(y)$ for a regular point y .

Choosing a generic homotopy $\{J_t\} \in \mathcal{J}_{\text{reg}}(W, \Omega; J_0, J_1)$ as provided by Proposition 5.11, the parametrized moduli space $\mathcal{M}_{0,1}^{A_0}(\{J_t\})$ now furnishes a smooth cobordism between the two closed manifolds $\mathcal{M}_{0,1}^{A_0}(J_0)$ and $\mathcal{M}_{0,1}^{A_0}(J_1)$.² Consider the evaluation map

$$\text{ev} : \mathcal{M}_{0,1}^{A_0}(\{J_t\}) \rightarrow W : ([S^2, j, z, u], t) \mapsto u(z),$$

and denote its restriction to the two boundary components by $\text{ev}_0 : \mathcal{M}_{0,1}^{A_0}(J_0) \rightarrow W$ and $\text{ev}_1 : \mathcal{M}_{0,1}^{A_0}(J_1) \rightarrow W$. As we saw in §5.1.1, ev_0 is a diffeomorphism, thus

²With a little more work, one can also give all of these moduli spaces natural orientations and thus obtain an *oriented* cobordism. This has the result that our use of the \mathbb{Z}_2 -mapping degree could be replaced by the integer-valued mapping degree, but we don't need this to prove the nonsqueezing theorem.

$(\text{ev}_0)_*[\mathcal{M}_{0,1}^{A_0}(J_0)] = [W] \in H_{2n}(W; \mathbb{Z}_2)$. It follows that

$$(\text{ev}_1)_*[\mathcal{M}_{0,1}^{A_0}(J_1)] = [W] \in H_{2n}(W; \mathbb{Z}_2)$$

as well, hence $\deg_2(\text{ev}_1) = 1$ and ev_1 is therefore surjective. In particular, $\text{ev}_1^{-1}(\iota(0))$ is not empty, and this proves Proposition 5.4.

5.2. Monotonicity in the integrable case

In this section, we consider only holomorphic curves in $\mathbb{R}^{2n} = \mathbb{C}^n$ with its standard complex structure i and symplectic structure ω_0 . Recall that a smooth map $u : \Sigma \rightarrow B_{r_0}$ is called **proper** if every compact set in the target has a compact preimage. For any $r \in (0, r_0)$, we define the compact subset

$$\Sigma_r := u^{-1}(\overline{B}_r^{2n}) \subset \Sigma,$$

which by Sard's theorem is a submanifold with smooth boundary for almost every r . Our main goal is to prove the following result, which was previously stated as Theorem 5.3 and was a crucial ingredient in the proof of the nonsqueezing theorem.

Theorem 5.12 (monotonicity). *if $u : (\Sigma, j) \rightarrow (B_{r_0}^{2n}, i)$ is a proper holomorphic map whose image contains 0, then for every $r \in (0, r_0)$,*

$$\int_{\Sigma_r} u^* \omega_0 \geq \pi r^2.$$

This result gives a quantitative version of the statement that a holomorphic curve cannot fit an arbitrarily small amount of area into some fixed neighborhood of a point in its image. More general versions also hold for non-integrable almost complex structures and are useful in proving a number of technical results, especially in the compactness theory; we'll come back to this in the next chapter. We should also mention that this kind of result is by no means unique to the theory of holomorphic curves: monotonicity formulas are also a popular tool in the theory of minimal surfaces (cf. [Law75, Grü88, CM99]), and indeed, Theorem 5.12 can be regarded as a corollary of such results after observing that whenever J is *compatible* with a symplectic structure ω and a Riemannian metric is defined by $\omega(\cdot, J\cdot)$, J -holomorphic curves are also *area minimizing*, cf. [MS04, Lemma 2.2.1]. This was also the perspective adopted by Gromov in [Gro85]; see also [Fis] for some more recent results along these lines. In order to keep the discussion self-contained and avoid delving into the theory of minimal surfaces, we shall instead present a direct “contact geometric” proof, which is fairly simple and uses a few notions that we will find useful in our later discussions of contact geometry.

To start with, it's easy to see from our knowledge of the local behavior of holomorphic curves that the estimate of Theorem 5.12 holds for any given curve u whenever $r > 0$ is sufficiently small. Indeed, in an appropriate choice of local coordinates on a small enough neighborhood, u looks like a small perturbation of the map

$$B_\epsilon \rightarrow \mathbb{C} \times \mathbb{C}^{n-1} : z \mapsto (z^k, 0),$$

whose area is $k\pi\epsilon^2$. (See §2.11 for a discussion of such local representation formulas.)

The result then follows from the next statement, which explains our use of the term “monotonicity”.

Proposition 5.13. *Given the setup of Theorem 5.12, the function*

$$F(r) = \frac{1}{r^2} \int_{\Sigma_r} u^* \omega_0$$

is nondecreasing.

Note that it will suffice to prove that $F(R) > F(r)$ whenever $0 < r < R < r_0$ and both r and R lie in the dense set of *regular* values, i.e. those for which the intersection of u with $\partial \overline{B}_r$ is transverse. For regular values, Σ_r is a smooth manifold with boundary and we can use Stokes’ theorem to compute $\int_{\Sigma_r} u^* \omega_0$. In order to uncover the dependence on r^2 , we shall switch perspectives and regard u as a map into the *symplectization of the standard contact sphere*.

Label the natural coordinates on $\mathbb{R}^{2n} = \mathbb{C}^n$ by $(z_1, \dots, z_n) = (p_1 + iq_1, \dots, p_n + iq_n)$, so the symplectic structure has the form

$$\omega_0 = \sum_{j=1}^n dp_j \wedge dq_j.$$

Recall from §1.6 that the vector field

$$Y_0 := \frac{1}{2} \sum_{j=1}^n \left(p_j \frac{\partial}{\partial p_j} + q_j \frac{\partial}{\partial q_j} \right)$$

is a *Liouville* vector field on $(\mathbb{R}^{2n}, \omega_0)$, meaning it satisfies $\mathcal{L}_{Y_0} \omega_0 = \omega_0$. Let λ_0 denote the 1-form on \mathbb{R}^{2n} which is ω -dual to Y_0 , i.e.

$$\lambda_0 := \omega_0(Y_0, \cdot).$$

An easy computation then produces the expression

$$\lambda_0 = \frac{1}{2} \sum_{j=1}^n (p_j dq_j - q_j dp_j),$$

and the fact that Y_0 is Liouville is equivalent to the observation that $d\lambda_0 = \omega_0$. Moreover, since $\lambda_0(Y_0) = \omega_0(Y_0, Y_0) = 0$, we also have

$$\mathcal{L}_{Y_0} \lambda_0 = \iota_{Y_0} d\lambda_0 + d\iota_{Y_0} \lambda_0 = \iota_{Y_0} \omega_0 = \lambda_0.$$

Identify the sphere S^{2n-1} with the boundary of the closed unit ball $\overline{B}^{2n} \subset \mathbb{R}^{2n}$, and define the **standard contact form** α_0 on S^{2n-1} as the restriction of λ_0 ,

$$\alpha_0 := \lambda_0|_{T(\partial \overline{B}^{2n})}.$$

Now consider the diffeomorphism

$$\Phi : \mathbb{R} \times S^{2n-1} \rightarrow \mathbb{R}^{2n} \setminus \{0\} : (t, m) \mapsto \varphi_{Y_0}^t(m) = e^{t/2} m,$$

where $\varphi_{Y_0}^t$ denotes the flow of Y_0 . By Exercise 1.33, we have

$$\Phi^* \lambda_0 = e^t \alpha_0, \quad \Phi^* \omega_0 = d(e^t \alpha_0),$$

where t denotes the \mathbb{R} -coordinate on $\mathbb{R} \times S^{2n-1}$ and α_0 is defined on $\mathbb{R} \times S^{2n-1}$ as the pullback via the projection $\mathbb{R} \times S^{2n-1} \rightarrow S^{2n-1}$. Define an integrable complex structure J_0 on $\mathbb{R} \times S^{2n-1}$ so that this diffeomorphism is biholomorphic, i.e.

$$J_0 := \Phi^* i.$$

Now removing at most finitely many points from Σ to define

$$\dot{\Sigma} := \{z \in \Sigma \mid u(z) \neq 0\}$$

and defining $\dot{\Sigma}_r \subset \Sigma_r$ similarly, we obtain a J_0 -holomorphic map

$$(u_{\mathbb{R}}, u_S) := \Phi^{-1} \circ u : \dot{\Sigma} \rightarrow \mathbb{R} \times S^{2n-1},$$

so that if $r = e^{\tau/2} \in (0, r_0)$ is regular, we have

$$\begin{aligned} F(r) &= \frac{1}{r^2} \int_{\Sigma_r} u^* \omega_0 = e^{-\tau} \int_{\dot{\Sigma}_r} (u_{\mathbb{R}}, u_S)^* d(e^t \alpha_0) = e^{-\tau} \int_{\partial \Sigma_r} (u_{\mathbb{R}}, u_S)^* (e^t \alpha_0) \\ &= \int_{\partial \Sigma_r} u_S^* \alpha_0. \end{aligned}$$

Thus for any two regular values $0 < r < R < r_0$, we now have

$$F(R) - F(r) = \int_{\partial \Sigma_R} u_S^* \alpha_0 - \int_{\partial \Sigma_r} u_S^* \alpha_0 = \int_{\Sigma_R \setminus \Sigma_r} u_S^* d\alpha_0.$$

Proposition 5.13 is then immediate from the following exercise.

Exercise 5.14. Show that the almost complex structure $J_0 = \Phi^* i$ on $\mathbb{R} \times S^{2n-1}$ has the following properties:

- (1) It is invariant under the natural \mathbb{R} -action by translation of the first factor in $\mathbb{R} \times S^{2n-1}$.
- (2) For any $t \in \mathbb{R}$, the unique hyperplane field in $\{t\} \times S^{2n-1}$ preserved by J_0 is precisely the *contact structure* $\xi_0 := \ker \alpha_0$.
- (3) The restriction of J_0 to ξ_0 is *compatible* with the symplectic bundle structure $d\alpha_0|_{\xi_0}$, i.e. the pairing $\langle X, Y \rangle := d\alpha_0(X, J_0 Y)$ defines a bundle metric on ξ_0 .
- (4) J_0 maps ∂_t to the *Reeb vector field* of α_0 , i.e. the unique vector field X_{α_0} on S^{2n-1} satisfying the conditions

$$d\alpha_0(X_{\alpha_0}, \cdot) \equiv 0 \quad \text{and} \quad \alpha_0(X_{\alpha_0}) \equiv 1.$$

Derive from these properties the fact that for any J_0 -holomorphic curve $(u_{\mathbb{R}}, u_S) : \Sigma \rightarrow \mathbb{R} \times S^{2n-1}$, the integrand $u_S^* d\alpha_0$ is nonnegative.

5.3. Bubbling off

Our goal in this section is to provide a *mostly* self-contained proof of Proposition 5.10, as a consequence of the following result.

Theorem 5.15. Suppose (M, ω) is a closed symplectic manifold of dimension $2n - 2 \geq 2$ with $\pi_2(M) = 0$, σ is an area form on S^2 , $W := S^2 \times M$, $\Omega := \sigma \oplus \omega$, $A_0 := [S^2 \times \{*\}] \in H_2(W)$ and we have the following sequences:

- $J_k \rightarrow J$ is a C^∞ -convergent sequence of Ω -compatible almost complex structures on W ,

- $u_k : (S^2, i) \rightarrow (W, J_k)$ is a sequence of pseudoholomorphic spheres with $[u_k] = A_0$, and
- $\zeta_k \in S^2$ is a sequence of marked points.

Then after taking a subsequence, there exist biholomorphic maps $\varphi_k : (S^2, i) \rightarrow (S^2, i)$ with $\varphi_k(0) = \zeta_k$ such that the reparametrized curves

$$u_k \circ \varphi_k : S^2 \rightarrow W$$

converge in C^∞ to a J -holomorphic sphere $u : (S^2, i) \rightarrow (W, J)$.

To prove this, we shall introduce some of the crucial technical tools that underlie the more general compactness results of the next chapter. There's only one result which we will need to take for now as a "black box":

Proposition 5.16 (Gromov's removable singularity theorem). *Suppose (M, ω) is a symplectic manifold with a tame almost complex structure J , and $u : B \setminus \{0\} \rightarrow M$ is a J -holomorphic curve which has finite energy $\int_{B \setminus \{0\}} u^* \omega < \infty$ and image contained in a compact subset of M . Then u extends smoothly over 0 to a J -holomorphic curve $B \rightarrow M$.*

A proof may be found in the next chapter, or in [MS04, Sik94, Hum97].

As a fundamental analytical tool for our compactness arguments, we will use the following piece of local elliptic regularity theory that was proved in Chapter 2 as Corollary 2.61:

Lemma 5.17. *Assume $p \in (2, \infty)$ and $m \geq 1$, $J_k \in \mathcal{J}^m(B^{2n})$ is a sequence of almost complex structures converging in C^m to $J \in \mathcal{J}^m(B^{2n})$, and $u_k : B \rightarrow B^{2n}$ is a sequence of J_k -holomorphic curves satisfying a uniform bound $\|u_k\|_{W^{1,p}(B)} < C$. Then u_k has a subsequence converging in $W_{\text{loc}}^{m+1,p}$ to a J -holomorphic curve $u : B \rightarrow B^{2n}$.*

In our situation, we have $J_k \rightarrow J$ in C^m for all m , thus we will obtain a C_{loc}^∞ -convergent subsequence if we can establish C^1 -bounds for our maps $u_k : S^2 \rightarrow W$, since C^1 embeds continuously into $W^{1,p}$. The lemma can be applied in a more global setting as follows. Fix Riemannian metrics on S^2 and W and use these to define the norm $|du(z)| \geq 0$ of the linear map $du(z) : T_z S^2 \rightarrow T_{u(z)} W$ for any $u \in C^1(S^2, W)$ and $z \in S^2$. If the given sequence of J_k -holomorphic maps $u_k : S^2 \rightarrow W$ satisfies a uniform bound of the form

$$(5.3) \quad |du_k(z)| < C \quad \text{for all } k \text{ and all } z \in S^2,$$

then since W is compact, a subsequence of u_k will converge in C^0 to some continuous map $u : S^2 \rightarrow W$. We can then cover both S^2 and $u(S^2) \subset W$ with finitely many local coordinate charts and apply Lemma 5.17, obtaining:

Lemma 5.18. *Suppose $J_k \rightarrow J$ is a C^∞ -convergent sequence of almost complex structures on a closed manifold W and $u_k : (S^2, i) \rightarrow (W, J_k)$ is a sequence of pseudoholomorphic curves satisfying a uniform C^1 -bound as in (5.3). Then a subsequence of u_k converges in C^∞ to a pseudoholomorphic curve $u : (S^2, i) \rightarrow (W, J)$.*

Remark 5.19. The above lemma is obviously also true if W is not compact but the images of the curves u_k are confined to a compact subset. This generalization is important for compactness results in contact geometry and symplectic field theory, e.g. [BEH⁺03].

In most situations, one cannot expect to derive a C^1 -bound directly from the given data, and in the general case such a bound does not even hold. The strategy is however as follows: if a C^1 -bound does *not* hold, then we can find a sequence of points $z_k \in S^2$ such that $|du_k(z_k)| \rightarrow \infty$, and by an intelligent choice of rescalings, the restriction of u_k to small neighborhoods of z_k gives rise to a sequence of holomorphic disks on expanding domains that exhaust \mathbb{C} . These disks are always nonconstant but satisfy a uniform C^1 -bound by construction, thus by Lemma 5.17 they will converge in C_{loc}^∞ to a J -holomorphic plane with finite energy. Since a plane is really just a punctured sphere, this J -holomorphic plane can be extended to a nonconstant holomorphic sphere, often called a “bubble”, and the process by which this sphere is extracted from the original sequence is often called “bubbling off”. In our situation, we will find that the existence of this bubble leads to a contradiction and thus implies the desired C^1 -bound on the original sequence. In more general settings, there is no contradiction and one must instead find a way of organizing the information that these bubbles add to the limit of the original sequence—this leads to the notion of *nodal* holomorphic curves, the more general objects that make up the *Gromov compactification*, to be discussed in the next chapter.

We now carry out the details of the above argument, using a particular type of rescaling trick that has been popularized by Hofer and collaborators (see e.g. [HZ94, §6.4]). The results stated below all assume the setting described in the statement of Theorem 5.15: in particular, $(W, \Omega) = (S^2 \times M, \sigma \oplus \omega)$ and $\pi_2(M) = 0$. Notice that the curves in the sequence $u_k : S^2 \rightarrow W$ are all homologous and thus all have the same energy

$$E(u_k) = \int_{S^2} u_k^* \Omega = \langle [\Omega], A_0 \rangle = \langle [\sigma], [S^2] \rangle = \int_{S^2} \sigma.$$

For reasons that will hopefully become clear in a moment, we now give this positive constant a special name and write

$$\hbar := \int_{S^2} \sigma > 0.$$

The following is then a very simple example of a general phenomenon known as *energy quantization*.

Lemma 5.20. *For any $J \in \mathcal{J}(W, \Omega)$, every nonconstant closed J -holomorphic sphere in W has energy at least \hbar .*

PROOF. If $u = (u_S, u_M) : S^2 \rightarrow S^2 \times M$ is J -holomorphic and not constant, then

$$\begin{aligned} 0 < E(u) &= \int_{S^2} u^* \Omega = \langle [\sigma \oplus \omega], [u_S] \times [\{*\}] + [\{*\}] \times [u_M] \rangle \\ &= \langle [\sigma], [u_S] \rangle + \langle [\omega], [u_M] \rangle. \end{aligned}$$

Since $\pi_2(M) = 0$, the spherical homology class $[u_M] \in H_2(M)$ necessarily vanishes, so the above expression implies $E(u) = \langle [\sigma], [u_S] \rangle$, which must be an integer multiple of \hbar . Since it is also positive, the result follows. \square

We next choose reparametrizations of the sequence u_k so as to rule out certain trivial possibilities, such as u_k converging almost everywhere to a constant. Write $u_k = (u_k^S, u_k^M) : S^2 \rightarrow S^2 \times M$, and observe that since $[u_k] = [S^2 \times \{*\}]$, $u_k^S : S^2 \rightarrow S^2$ is always a map of degree 1 and hence surjective. After taking a subsequence, we may assume that the images of the marked points in S^2 converge, i.e.

$$u_k^S(\zeta_k) \rightarrow \zeta_\infty \in S^2.$$

Assume without loss of generality that ζ_∞ is neither 1 nor ∞ ; if it is one of these, then the remainder of our argument will require only trivial modifications. Now since u_k^S is surjective, for sufficiently large k we can always find biholomorphic maps $\varphi_k : (S^2, i) \rightarrow (S^2, i)$ that have the following properties:

- $\varphi_k(0) = \zeta_k$,
- $u_k^S \circ \varphi_k(1) = 1$,
- $u_k^S \circ \varphi_k(\infty) = \infty$.

To simplify notation, let us now replace the original sequence by these reparametrizations and thus assume without loss of generality that the maps $u_k = (u_k^S, u_k^M) : S^2 \rightarrow S^2 \times M$ and marked points $\zeta_k \in S^2$ satisfy

$$\zeta_k = 0, \quad u_k(1) \in \{1\} \times M, \quad u_k(\infty) \in \{\infty\} \times M$$

for all k .

If the maps u_k satisfy a uniform C^1 -bound, then we are now finished due to Lemma 5.18. Thus assume the contrary, that there is a sequence $z_k \in S^2$ with

$$|du_k(z_k)| \rightarrow \infty,$$

and after taking a subsequence we may assume $z_k \rightarrow z_\infty \in S^2$. Choose a neighborhood $z_\infty \in \mathcal{U} \subset S^2$ and a biholomorphic map

$$\varphi : (B, i) \rightarrow (\mathcal{U}, i)$$

identifying \mathcal{U} with the unit ball in \mathbb{C} such that $\varphi(0) = z_\infty$, and write

$$\tilde{u}_k = u_k \circ \varphi : (B, i) \rightarrow (W, J_k), \quad \tilde{z}_k = \varphi^{-1}(z_k).$$

We then have $|d\tilde{u}_k(\tilde{z}_k)| \rightarrow \infty$ and $\tilde{z}_k \rightarrow 0$.

We now examine a rescaled reparametrization of the sequence \tilde{u}_k on shrinking neighborhoods of \tilde{z}_k . In particular, let $R_k := |d\tilde{u}_k(\tilde{z}_k)| \rightarrow \infty$, pick a sequence of positive numbers $\epsilon_k \rightarrow 0$ which decay slowly enough so that $\epsilon_k R_k \rightarrow \infty$, and consider the sequence of J_k -holomorphic maps

$$v_k : (B_{\epsilon_k R_k}, i) \rightarrow (W, J_k) : z \mapsto \tilde{u}_k \left(\tilde{z}_k + \frac{z}{R_k} \right).$$

Then

$$|dv_k(z)| = \frac{1}{R_k} \left| d\tilde{u}_k \left(\tilde{z}_k + \frac{z}{R_k} \right) \right|,$$

so in particular $|dv_k(0)| = \frac{1}{R_k} |d\tilde{u}_k(\tilde{z}_k)| = 1$. To proceed further, we'd like to be able to say that $|dv_k(z)|$ satisfies a uniform bound for $z \in B_{\epsilon_k R_k}$, as then Lemma 5.17 would give a subsequence converging in C_{loc}^∞ on \mathbb{C} . Such a bound is not obvious: it would require being able to bound $|d\tilde{u}_k(z)|$ in terms of $|d\tilde{u}_k(\tilde{z}_k)|$ for all $z \in B_{\epsilon_k}(\tilde{z}_k)$. While there is no reason that such a bound should necessarily hold for the chosen sequence, the following topological lemma due to Hofer tells us that we can always ensure this bound after a slight adjustment.

Lemma 5.21 (Hofer). *Suppose (X, d) is a complete metric space, $g : X \rightarrow [0, \infty)$ is continuous, $x_0 \in X$ and $\epsilon_0 > 0$. Then there exist $x \in X$ and $\epsilon > 0$ such that,*

- (a) $\epsilon \leq \epsilon_0$,
- (b) $g(x)\epsilon \geq g(x_0)\epsilon_0$,
- (c) $d(x, x_0) \leq 2\epsilon_0$, and
- (d) $g(y) \leq 2g(x)$ for all $y \in \overline{B_\epsilon(x)}$.

PROOF. If there is no $x_1 \in \overline{B_{\epsilon_0}(x_0)}$ such that $g(x_1) > 2g(x_0)$, then we can set $x = x_0$ and $\epsilon = \epsilon_0$ and are done. If such a point x_1 does exist, then we set $\epsilon_1 := \epsilon_0/2$ and repeat the above process for the pair (x_1, ϵ_1) : that is, if there is no $x_2 \in \overline{B_{\epsilon_1}(x_1)}$ with $g(x_2) > 2g(x_1)$, we set $(x, \epsilon) = (x_1, \epsilon_1)$ and are finished, and otherwise define $\epsilon_2 = \epsilon_1/2$ and repeat for (x_2, ϵ_2) . This process must eventually terminate, as otherwise we obtain a Cauchy sequence x_n with $g(x_n) \rightarrow \infty$, which is impossible if X is complete. \square

The upshot of the lemma is that the sequences $\epsilon_k > 0$ and $\tilde{z}_k \in B$ can be modified slightly to have the additional property that

$$(5.4) \quad |d\tilde{u}_k(z)| \leq 2|d\tilde{u}_k(\tilde{z}_k)| \quad \text{for all } z \in \overline{B_{\epsilon_k}(\tilde{z}_k)}.$$

From this it follows that the rescaled sequence $v_k : B_{\epsilon_k R_k} \rightarrow W$ satisfies

$$|dv_k(z)| \leq 2, \quad |dv_k(0)| = 1,$$

so we conclude from Lemma 5.17 that a subsequence of v_k converges in $C_{\text{loc}}^\infty(\mathbb{C}, W)$ to a J -holomorphic plane

$$v_\infty : (\mathbb{C}, i) \rightarrow (W, J)$$

which satisfies $|dv_\infty(0)| = 1$ and is thus not constant. We claim that v_∞ also has finite energy bounded by \hbar . Indeed, for any $R > 0$, we have

$$\int_{B_R} v_\infty^* \Omega = \lim_k \int_{B_R} v_k^* \Omega,$$

while for sufficiently large k ,

$$\int_{B_R} v_k^* \Omega \leq \int_{B_{\epsilon_k R_k}} v_k^* \Omega = \int_{B_{\epsilon_k}(\tilde{z}_k)} \tilde{u}_k^* \Omega = \int_{\varphi(B_{\epsilon_k}(\tilde{z}_k))} u_k^* \Omega \leq \int_{S^2} u_k^* \Omega = \hbar.$$

Applying the removable singularity theorem (Prop. 5.16) and energy quantization (Lemma 5.20), v_∞ thus extends to a nonconstant J -holomorphic sphere

$$v_\infty : (S^2, i) \rightarrow (W, J)$$

with energy exactly \hbar . This sphere is our first real life example of a so-called “bubble”.

We claim next that if the above scenario happens, then for any other sequence $z'_k \in S^2$ with $|du_k(z'_k)| \rightarrow \infty$, z'_k can only accumulate at the same point z_∞ again. Indeed, otherwise the above procedure produces a second bubble $v'_\infty : (S^2, i) \rightarrow (W, J)$ with energy \hbar , and by inspecting the energy estimate above, one sees that for large k , u_k must have a concentration of energy close to \hbar in small neighborhoods of both z_∞ and z'_∞ . That is impossible since $E(u_k)$ is already bounded by \hbar .

The above implies that on any compact subset of $S^2 \setminus \{z_\infty\}$, u_k satisfies a uniform C^1 -bound and thus converges in $C_{\text{loc}}^\infty(S^2 \setminus \{z_\infty\})$ to a J -holomorphic punctured sphere

$$u_\infty : (S^2 \setminus \{z_\infty\}, i) \rightarrow (W, J).$$

Moreover, we have

$$u_\infty(0) \in \{\zeta_\infty\} \times M, \quad u_\infty(1) \in \{1\} \times M \quad \text{and} \quad u_\infty(\infty) \in \{\infty\} \times M$$

unless $z_\infty \in \{\zeta_\infty, 1, \infty\}$, in which case at least two of these three statements still holds. It follows that u_∞ cannot be constant, so by Lemma 5.20 it has energy at least \hbar . But this again gives a contradiction if the bubble v_∞ exists, as it implies that for large k , the restrictions of u_k to some large subset of $S^2 \setminus \{z_\infty\}$ and some disjoint small neighborhood of z_∞ each have energy at least slightly less than \hbar , so that $\int_{S^2} u_k^* \Omega$ must be strictly greater than \hbar . This contradiction excludes the bubbling scenario, thus establishing the desired C^1 -bound for u_k and completing the proof of Theorem 5.15.

Bibliography

- [AH] C. Abbas and H. Hofer, *Holomorphic curves and global questions in contact geometry*. In preparation.
- [AA02] Y. A. Abramovich and C. D. Aliprantis, *An invitation to operator theory*, Graduate Studies in Mathematics, vol. 50, American Mathematical Society, Providence, RI, 2002.
- [ADN59] S. Agmon, A. Douglis, and L. Nirenberg, *Estimates near the boundary for solutions of elliptic partial differential equations satisfying general boundary conditions. I*, Comm. Pure Appl. Math. **12** (1959), 623–727.
- [ADN64] ———, *Estimates near the boundary for solutions of elliptic partial differential equations satisfying general boundary conditions. II*, Comm. Pure Appl. Math. **17** (1964), 35–92.
- [Arn89] V. I. Arnol’d, *Mathematical methods of classical mechanics*, Springer-Verlag, New York, 1989.
- [Aud94] M. Audin, *Symplectic and almost complex manifolds*, Holomorphic curves in symplectic geometry, Progr. Math., vol. 117, Birkhäuser, Basel, 1994, pp. 41–74. With an appendix by P. Gauduchon.
- [ALP94] M. Audin, F. Lalonde, and L. Polterovich, *Symplectic rigidity: Lagrangian submanifolds*, Holomorphic curves in symplectic geometry, Progr. Math., vol. 117, Birkhäuser, Basel, 1994, pp. 271–321.
- [BEH⁺03] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder, *Compactness results in symplectic field theory*, Geom. Topol. **7** (2003), 799–888 (electronic).
- [CZ52] A. P. Calderon and A. Zygmund, *On the existence of certain singular integrals*, Acta Math. **88** (1952), 85–139.
- [CZ56] A. P. Calderón and A. Zygmund, *On singular integrals*, Amer. J. Math. **78** (1956), 289–309.
- [CE] K. Cieliebak and Y. Eliashberg, *Symplectic Geometry of Stein Manifolds*. In preparation.
- [CM99] T. H. Colding and W. P. Minicozzi II, *Minimal surfaces*, Courant Lecture Notes in Mathematics, vol. 4, New York University Courant Institute of Mathematical Sciences, New York, 1999.
- [DK90] S. K. Donaldson and P. B. Kronheimer, *The geometry of four-manifolds*, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 1990. Oxford Science Publications.
- [DN55] A. Douglis and L. Nirenberg, *Interior estimates for elliptic systems of partial differential equations*, Comm. Pure Appl. Math. **8** (1955), 503–538.
- [Eel66] J. Eells Jr., *A setting for global analysis*, Bull. Amer. Math. Soc. **72** (1966), 751–807.
- [Eli89] Y. Eliashberg, *Classification of overtwisted contact structures on 3-manifolds*, Invent. Math. **98** (1989), no. 3, 623–637.
- [Eli90] Y. Eliashberg, *Filling by holomorphic discs and its applications*, Geometry of low-dimensional manifolds, 2 (Durham, 1989), London Math. Soc. Lecture Note Ser., vol. 151, Cambridge Univ. Press, Cambridge, 1990, pp. 45–67.
- [Eli96] Y. Eliashberg, *Unique holomorphically fillable contact structure on the 3-torus*, Internat. Math. Res. Notices **2** (1996), 77–82.
- [Eli67] H. I. Eliasson, *Geometry of manifolds of maps*, J. Differential Geometry **1** (1967), 169–194.

- [Eva98] L. C. Evans, *Partial differential equations*, Graduate Studies in Mathematics, vol. 19, American Mathematical Society, Providence, RI, 1998. MR1625845 (99e:35001)
- [FK92] H. M. Farkas and I. Kra, *Riemann surfaces*, 2nd ed., Graduate Texts in Mathematics, vol. 71, Springer-Verlag, New York, 1992.
- [Fis] J. W. Fish, *Estimates for J-curves as submanifolds*. Preprint arXiv:0912.4445, to appear in Geom. Topol.
- [Flo88] A. Floer, *The unregularized gradient flow of the symplectic action*, Comm. Pure Appl. Math. **41** (1988), no. 6, 775–813.
- [FO99] K. Fukaya and K. Ono, *Arnold conjecture and Gromov-Witten invariant*, Topology **38** (1999), no. 5, 933–1048.
- [Gay06] D. T. Gay, *Four-dimensional symplectic cobordisms containing three-handles*, Geom. Topol. **10** (2006), 1749–1759 (electronic).
- [Gei06] H. Geiges, *Contact geometry*, Handbook of differential geometry. Vol. II, Elsevier/North-Holland, Amsterdam, 2006, pp. 315–382.
- [Gei08] ———, *An introduction to contact topology*, Cambridge Studies in Advanced Mathematics, vol. 109, Cambridge University Press, Cambridge, 2008.
- [GH] P. Ghiggini and K. Honda, *Giroux torsion and twisted coefficients*. Preprint arXiv:0804.1568.
- [Gir94] E. Giroux, *Une structure de contact, même tendue, est plus ou moins tordue*, Ann. Sci. École Norm. Sup. (4) **27** (1994), no. 6, 697–705 (French, with English summary).
- [Gro85] M. Gromov, *Pseudoholomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), no. 2, 307–347.
- [Grü88] M. Grüter, *The monotonicity formula in geometric measure theory, and an application to a partially free boundary problem*, Partial differential equations and calculus of variations, Lecture Notes in Math., vol. 1357, Springer, Berlin, 1988, pp. 238–255.
- [Hof93] H. Hofer, *Pseudoholomorphic curves in symplectizations with applications to the Weinstein conjecture in dimension three*, Invent. Math. **114** (1993), no. 3, 515–563.
- [Hof] ———, *Polyfolds and a general Fredholm theory*. Preprint arXiv:0809.3753.
- [HZ90] H. Hofer and E. Zehnder, *A new capacity for symplectic manifolds*, Analysis, et cetera, Academic Press, Boston, MA, 1990, pp. 405–427.
- [HZ94] H. Hofer and E. Zehnder, *Symplectic invariants and Hamiltonian dynamics*, Birkhäuser Verlag, Basel, 1994.
- [Hum97] C. Hummel, *Gromov’s compactness theorem for pseudo-holomorphic curves*, Progress in Mathematics, vol. 151, Birkhäuser Verlag, Basel, 1997.
- [IS99] S. Ivashkovich and V. Shevchishin, *Structure of the moduli space in a neighborhood of a cusp-curve and meromorphic hulls*, Invent. Math. **136** (1999), no. 3, 571–602.
- [Lan93] S. Lang, *Real and functional analysis*, 3rd ed., Springer-Verlag, New York, 1993.
- [Lan99] ———, *Fundamentals of differential geometry*, Springer-Verlag, New York, 1999.
- [Law75] H. B. Lawson Jr., *Minimal varieties*, Differential geometry (Proc. Sympos. Pure Math., Vol. XXVII, Stanford Univ., Stanford, Calif., 1973), Part 1, Amer. Math. Soc., Providence, R.I., 1975, pp. 143–175.
- [McD90] D. McDuff, *The structure of rational and ruled symplectic 4-manifolds*, J. Amer. Math. Soc. **3** (1990), no. 3, 679–712.
- [MS98] D. McDuff and D. Salamon, *Introduction to symplectic topology*, The Clarendon Press Oxford University Press, New York, 1998.
- [MS04] ———, *J-holomorphic curves and symplectic topology*, American Mathematical Society, Providence, RI, 2004.
- [MW95] M. J. Micallef and B. White, *The structure of branch points in minimal surfaces and in pseudoholomorphic curves*, Ann. of Math. (2) **141** (1995), no. 1, 35–85.
- [Mil97] J. W. Milnor, *Topology from the differentiable viewpoint*, Princeton Landmarks in Mathematics, Princeton University Press, Princeton, NJ, 1997. Based on notes by David W. Weaver; Revised reprint of the 1965 original.

- [OS04] B. Ozbagci and A. I. Stipsicz, *Surgery on contact 3-manifolds and Stein surfaces*, Bolyai Society Mathematical Studies, vol. 13, Springer-Verlag, Berlin, 2004.
- [Pal68] R. S. Palais, *Foundations of global non-linear analysis*, W. A. Benjamin, Inc., New York-Amsterdam, 1968.
- [Rab78] P. H. Rabinowitz, *Periodic solutions of Hamiltonian systems*, Comm. Pure Appl. Math. **31** (1978), no. 2, 157–184.
- [SS92] M. Seppälä and T. Sorvali, *Geometry of Riemann surfaces and Teichmüller spaces*, North-Holland Mathematics Studies, vol. 169, North-Holland Publishing Co., Amsterdam, 1992.
- [Sie08] R. Siefring, *Relative asymptotic behavior of pseudoholomorphic half-cylinders*, Comm. Pure Appl. Math. **61** (2008), no. 12, 1631–1684.
- [Sik94] J.-C. Sikorav, *Some properties of holomorphic curves in almost complex manifolds*, Holomorphic curves in symplectic geometry, Progr. Math., vol. 117, Birkhäuser, Basel, 1994, pp. 165–189.
- [Str90] M. Struwe, *Existence of periodic solutions of Hamiltonian systems on almost every energy surface*, Bol. Soc. Brasil. Mat. (N.S.) **20** (1990), no. 2, 49–58.
- [Tau07] C. H. Taubes, *The Seiberg-Witten equations and the Weinstein conjecture*, Geom. Topol. **11** (2007), 2117–2202.
- [Tay96] M. E. Taylor, *Partial differential equations. I*, Applied Mathematical Sciences, vol. 115, Springer-Verlag, New York, 1996. Basic theory.
- [Tro92] A. J. Tromba, *Teichmüller theory in Riemannian geometry*, Lectures in Mathematics ETH Zürich, Birkhäuser Verlag, Basel, 1992. Lecture notes prepared by Jochen Denzler.
- [Vek62] I. N. Vekua, *Generalized analytic functions*, Pergamon Press, London, 1962.
- [Vit87] C. Viterbo, *A proof of Weinstein’s conjecture in \mathbf{R}^{2n}* , Ann. Inst. H. Poincaré Anal. Non Linéaire **4** (1987), no. 4, 337–356 (English, with French summary).
- [Wei78] A. Weinstein, *Periodic orbits for convex Hamiltonian systems*, Ann. of Math. (2) **108** (1978), no. 3, 507–518.
- [Wen10a] C. Wendl, *Automatic transversality and orbifolds of punctured holomorphic curves in dimension four*, Comment. Math. Helv. **85** (2010), no. 2, 347–407.
- [Wen10b] ———, *Strongly fillable contact manifolds and J-holomorphic foliations*, Duke Math. J. **151** (2010), no. 3, 337–384.